ANIMA WG
Internet-Draft
Intended status: Informational
Expires: July 14, 2016

S. Jiang, Ed.
Z. Du
Huawei Technologies Co., Ltd
B. Carpenter
Univ. of Auckland
Q. Sun
China Telecom

January 11, 2016

Autonomic IPv6 Edge Prefix Management in Large-scale Networks
draft-ietf-anima-prefix-management-00

Abstract

This document describes an autonomic solution for IPv6 prefix management at the edge of large-scale ISP networks. An important purpose of the document is to use it for validation of the design of various components of the autonomic networking infrastructure.

Table of Contents

1. Introduction

This document proposes an autonomic solution for IPv6 prefix management in large-scale networks. The background to Autonomic Networking (AN) is described in [RFC7575] and [RFC7576]. A generic autonomic signaling protocol (GRASP) is specified by [I-D.ietf-anima-grasp] and would be used by the proposed autonomic prefix management solution. An important purpose of the present document is to use it for validation of the design of GRASP and other components of the autonomic networking infrastructure described in [I-D.behringer-anima-reference-model].

This document is not intended to solve all cases of IPv6 prefix management. In fact, it assumes that the network's main infrastructure elements already have addresses and prefixes. The document is dedicated to how to make IPv6 prefix management at the edges of large-scale networks as autonomic as possible. It is specifically written for service provider (ISP) networks. Although there are similarities between ISPs and large enterprise networks, the requirements for the two use cases differ.

However, the solution is designed in a general way. Its use for a broader scope than edge prefixes, including some or all infrastructure prefixes, is left for future discussion.

Note in draft: This version is preliminary. In particular, many design details may be subject to change until the Anima specifications become agreed.

2. Terminology

TBD

3. Problem Statement

The autonomic networking use case considered here is autonomic IPv6 prefix management at the edge of large-scale ISP networks.

Although DHCPv6 Prefix Delegation [RFC3633] supports automated delegation of IPv6 prefixes from one router to another, prefix management is still largely depending on human planning. In other words, there is no basic information or policy to support autonomic decisions on the prefix length that each router should request or be delegated, according to its role in the network. Roles could be locally defined or could be generic (edge router, interior router, etc.). Furthermore, IPv6 prefix management by humans tends to be rigid and static after initial planning.

The problem to be solved by autonomic networking is how to dynamically manage IPv6 address space in large-scale networks, so that IPv6 addresses can be used efficiently. Here, we limit the problem to assignment of prefixes at the edge of the network, close to access routers that support individual fixed-line subscribers, mobile customers, and corporate customers. We assume that the core infrastructure of the network has already been established with appropriately assigned prefixes. The AN approach discussed in this document is based on the assumption that there is a generic discovery and negotiation protocol that enables direct negotiation between intelligent IP routers. GRASP [I-D.ietf-anima-grasp] is intended to be such a protocol.

3.1. Intended User and Administrator Experience

The intended experience is, for the administrator(s) of a large-scale network, that the management of IPv6 address space at the edge of the network can be run with minimum efforts, as devices at the edge are added and removed and as customers of all kinds join and leave the network. In the ideal scenario, the administrator(s) only have to specify a single IPv6 prefix for the whole network and the initial prefix length for each device role. As far as users are concerned, IPv6 prefix assignment would occur exactly as it does in any other network.

The actual prefix usage needs to be logged for potential offline management operations including audit and security incident tracing.

3.2. Analysis of Parameters and Information Involved

For specific purposes of address management, a few parameters are involved on each edge device (some of them can be pre-configured before they are connected). They include:

- o Identity, authentication and authorization of this device. This is expected to use the autonomic networking secure bootstrap process [I-D.ietf-anima-bootstrapping-keyinfra], following which the device could safely take part in autonomic operations.

- o Role of this device.

- o An IPv6 prefix length for this device.

- o An IPv6 prefix that is assigned to this device and its downstream devices.

A few parameters are involved in the network as a whole. They are:

- o Identity of a trust anchor, which is a certification authority (CA) maintained by the network administrator(s), used during the secure bootstrap process.

- o Total IPv6 address space available for edge devices. It is one (or several) IPv6 prefix(es).

- o The initial prefix length for each device role.

3.2.1. Parameters each device can decide for itself

This section identifies those of the above parameters that do not need external information in order for the devices concerned to set them to a reasonable value after bootstrap or after a network disruption. There are few of these:

- Role of this device.
- Default IPv6 prefix length for this device.
- Identity of this device.

The device may be shipped from the manufacturer with pre-configured role and default prefix length, which could be modified by an autonomic mechanism.

3.2.2. Information needed from policy intent

This section identifies those parameters that need external information about policy intent in order for the devices concerned to set them to a non-default value.

- Non-default value for the IPv6 prefix length for this device. This needs to be decided based on the role of this device.
- The initial prefix length for each device role.
- Whether to allow the device request more address space.
- The policy when to request more address space, for example, if the address usage reaches a certain limit or percentage.

3.2.3. Comparison with current solutions

This section briefly compares the above use case with current solutions. Currently, the address management is still largely dependent on human planning. It is rigid and static after initial planning. Address requests will fail if the configured address space is used up.

Some autonomic and dynamic address management functions may be achievable by extending the existing protocols, for example, extending DHCPv6-PD to request IPv6 prefixes according to the device role. However, defining uniform device roles may not be a practical task. Some functions are not suitable to be achieved by any existing protocols.

Using a generic autonomic discovery and negotiation protocol instead of specific solutions has the advantage that additional parameters can be included in the autonomic solution without creating new mechanisms. This is the principal argument for a generic approach.

3.3. Interaction with other devices

3.3.1. Information needed from other devices

This section identifies those of the above parameters that need external information from neighbor devices (including the upstream devices). In many cases, two-way dialogue with neighbor devices is needed to set or optimize them.

- Identity of a trust anchor.
- The device will need to discover a device, from which it can acquire IPv6 address space.
- The initial prefix length for each device role, particularly for its own downstream devices.
- The default value of the IPv6 prefix length may be overridden by a non-default value.
- The device will need to request and acquire IPv6 prefix that is assigned to this device and its downstream devices.
- The device may respond to prefix delegation request from its downstream devices.
- The device may require to be assigned more IPv6 address space, if it used up its assigned IPv6 address space.

3.3.2. Monitoring, diagnostics and reporting

This section discusses what role devices should play in monitoring, fault diagnosis, and reporting.

- The actual address assignments need to be logged for the potential offline management operations.
- In general, the usage situation of address space should be reported to the network administrators, in an abstract way, for example, statistics or visualized report.
- A forecast of address exhaustion should be reported.

4. Autonomic Edge Prefix Management Solution

This section introduces an autonomic edge prefix management solution. It uses the generic discovery and negotiation protocol defined by [I-D.ietf-anima-grasp]. The relevant options are defined in Section 5.

The procedures described below are carried out by an Autonomic Service Agent (ASA) in each device that participates in the solution. We will refer to this as the PrefixManager ASA.

4.1. Behaviors on prefix requesting device

If the device containing an PrefixManager ASA has used up its address pool, it can request more space according to its requirements. It should decide the length of the requested prefix by the intent-based mechanism, described in Section 6.

An PrefixManager ASA that needs additional address space should firstly discover peers that may be able to provide extra address space. The ASA should send out a GRASP Discovery message that contains an PrefixManager Objective option Section 5.1 in order to discover peers also supporting that option. Then it should choose one such peer, most likely the first to respond.

If the GRASP discovery Response message carries a divert option pointing to an off-link PrefixManager ASA, the requesting ASA may initiate negotiation with that ASA diverted device to find out whether it can provide the requested length prefix.

In any case, the requesting ASA will act as a GRASP negotiation initiator by sending a GRASP Request message with an PrefixManager Objective option. The ASA indicates in this option both the length of the requested prefix and whether the ASA supports the DHCPv6 Prefix Delegation (PD) function [RFC3633]. This starts a GRASP negotiation process.

During the subsequent negotiation, the ASA will decide at each step whether to accept the offered prefix. That decision, and the decision to end negotiation, is an implementation choice.

The ASA could alternatively initiate rapid mode GRASP discovery with an embedded negotiation request, if it is implemented.

4.2. Behaviors on prefix providing device

A device that receives a Discovery message with an PrefixManager Objective option should respond with a GRASP Response message if it contains an PrefixManager ASA. Further details of the discovery process are described in [I-D.ietf-anima-grasp]. When this ASA receives a subsequent Request message it should conduct a GRASP negotiation sequence, using Negotiate, Confirm-waiting, and Negotiation-ending messages as appropriate. The Negotiate messages carry an PrefixManager Objective option. This will indicate whether the sending device supports the PD function. More importantly, it will indicate the prefix and its length offered to the requesting ASA. As described in [I-D.ietf-anima-grasp], negotiation will continue until either end stops it with a Negotiation-ending message. If the negotiation succeeds, the prefix providing ASA will remove the negotiated prefix from its pool, and the requesting ASA will add it. If the negotiation fails, the party sending the Negotiation-ending message may include an error code string.

During the negotiation, the ASA will decide at each step how large a prefix to offer. That decision, and the decision to end negotiation, is an implementation choice.

The ASA could alternatively negotiate in response to rapid mode GRASP discovery, if it is implemented.

This specification is independent of whether the PrefixManager ASAs are all embedded in routers, but that would be a rather natural scenario. A gateway router in a hierarchical network topology normally provides prefixes for routers within its subnet, and it is likely to contain the first PrefixManager ASA discovered by its downstream routers. However, the GRASP discovery model, including its Redirect feature, means that this is not an exclusive scenario, and a downstream PrefixManager ASA could negotiate a new prefix with a router other than its upstream router.

A resource shortage may cause the gateway router to request more resource in turn from its own upstream device. This would be another independent GRASP discovery and negotiation process. During the processing time, the gateway router should send a Confirm-waiting Message to the initial requesting router, to extend its timeout. When the new resource becomes available, the gateway router responds with a GRASP Negotiate message with a prefix length matching the request.

The algorithm to choose which prefixes to assign on the prefix providing devices is an implementation choice.

4.3. Behavior after Successful Negotiation

Upon receiving a GRASP Negotiation-ending message that indicates that an acceptable prefix length is available, the requesting device may request the prefix using DHCPv6 PD, if both ASAs have indicated that they are within a device that supports PD. Otherwise, it is permissible for the initiating ASA to use the negotiated prefix without further messages.

[Author's note: It is not intended to undermine DHCPv6 PD. But in fact, if PD is not supported and the GRASP negotiation has succeeded, there should be no problem with this and it seems consistent as a solution.]

4.4. Prefix logging

Within the autonomic prefix management, all the prefix assignment is done by devices without human intervention. It is therefore important to record all the prefix assignment history. However, the logging and reporting process is out of scope for this specification.

5. Autonomic Prefix Management Options

This section defines the GRASP options that are used to support autonomic prefix management.

5.1. Edge Prefix Objective Option

The PrefixManager Objective option is a GRASP objective option conforming to [I-D.ietf-anima-grasp]. Its name is "PrefixManager" (see Section 8) and it carries up to three data items as its value: the PD support flag, the prefix length, and the actual prefix bits. The format of the PrefixManager Objective option is described as follows in CBOR data definition language (CDDL) [I-D.greevenbosch-appsawg-cbor-cddl]:

```
objective = ["PrefixManager", objective-flags, loop-count,
             PD-support, length, ?prefix]

loop-count = 0..255         ; as in the GRASP specification
objective-flags /=          ; as in the GRASP specification
PD-support = true / false   ; indicates whether sender supports PD
length = 0..128             ; requested or offered prefix length
prefix = bytes .size 16     ; offered prefix in binary format
```

6. Prefix Management Intent

With in a single administrative domain, the network operator could provide intent for all devices with a certain role. Thus it would be possible to apply an intended policy for every device in a simple way, without human intervention or configuration files.

For example, the network operator could define the default prefix length for each type of role. A prefix management intent, which contains all mapping information of device roles and their default prefix lengths, should be flooded in the network, through the Autonomic Control Plane (ACP) [I-D.ietf-anima-autonomic-control-plane]. The intent flooding mechanism is not yet defined, but one possibility would be define a suitable GRASP synchronization objective and flood it through the network. To make this concrete, there could be an objective defined as follows:

```
objective = ["Intent.PrefixManager", objective-flags, text]

loop-count = 0..255         ; as in the GRASP specification
objective-flags /=          ; as in the GRASP specification

;The text object would be the relevant intent statements (such
;as the example below) transmitted as a single string with all
;whitespace and format characters removed.
```

This could be flooded to all nodes, and any PrefixManager ASA that did not receive it for some reason could obtain a copy using GRASP synchronization. Upon receiving the prefix management intent, every device can decide its default prefix length by matching its own role.

6.1. Example of Prefix Management Intent

The prefix management intent in this document is used to carry mapping information of device roles and their default prefix lengths in an autonomic domain. For example, an IPRAN operator wants to configure the prefix length of RNC Site Gateway (RSG) as 34, the prefix length of Aggregation Site Gateway (ASG) as 44, and the prefix length of Cell Site Gateway (CSG) as 56. She/he may input the following intent into the autonomic network:

```
{"autonomic_intent":
[
   {"model_version": "1.0"},
   {"intent_type": "Network management"},
   {"autonomic_domain": "Customer_X_intranet"},
   {"intent_name": "Prefix management"},
   {"intent_version": 73},
   {"Timestamp": "20150606 00:00:00"},
   {"Lifetime": "Permanent"},
   {"signature": "XXXXXXXXXXXXXXXXXXXX"},
   {"content":
   [
      {"role": [{"role_name": "RSG"},
         {"role_characteristic":
            [{"prefix_length": "34"}]}
         ]},
      {"role": [{"role_name": "ASG"},
         {"role_characteristic":
            [{"prefix_length": "44"}]}
         ]},
      {"role": [{"role_name": "CSG"},
         {"role_characteristic":
            [{"prefix_length": "56"}]}
         ]}
   ]
   }
]
}
```


7. Security Considerations

Relevant security issues are discussed in [I-D.ietf-anima-grasp]. The preferred security model is that devices are trusted following the secure bootstrap procedure [I-D.ietf-anima-bootstrapping-keyinfra] and that a secure Autonomic Control Plane (ACP) [I-D.ietf-anima-autonomic-control-plane] is in place.

It is RECOMMENDED that DHCPv6 PD, if used, should be operated using DHCPv6 authentication or Secure DHCPv6.

8. IANA Considerations

This document defines one new GRASP Objective Option name, "PrefixManager". The IANA is requested to add this to the GRASP Objective Names Table registry defined by [I-D.ietf-anima-grasp] (if approved).

9. Acknowledgements

Valuable comments were received from Michael Behringer, Joel Halpern, and Chongfeng Xie.

This document was produced using the xml2rfc tool [RFC2629].

10. Change log [RFC Editor: Please remove]

draft-jiang-anima-prefix-management-00: original version, 2014-10-25.

draft-jiang-anima-prefix-management-01: add intent example and coauthor Zongpeng Du, 2015-05-04.

draft-jiang-anima-prefix-management-02: update references and the format of the prefix management intent, 2015-10-14.

draft-ietf-anima-prefix-management-00: WG adoption, clarify scope and purpose, update text to match latest GRASP spec, 2016-01-11.

11. References

11.1. Normative References

[I-D.greevenbosch-appsawg-cbor-cddl]
Vigano, C. and H. Birkholz, "CBOR data definition language (CDDL): a notational convention to express CBOR data structures", draft-greevenbosch-appsawg-cbor-cddl-07 (work in progress), October 2015.

[I-D.ietf-anima-autonomic-control-plane]
Behringer, M., Bjarnason, S., BL, B., and T. Eckert, "An Autonomic Control Plane", draft-ietf-anima-autonomic-control-plane-01 (work in progress), October 2015.

[I-D.ietf-anima-bootstrapping-keyinfra]
Pritikin, M., Richardson, M., Behringer, M., and S. Bjarnason, "Bootstrapping Key Infrastructures", draft-ietf-anima-bootstrapping-keyinfra-01 (work in progress), October 2015.

[I-D.ietf-anima-grasp]
Bormann, C., Carpenter, B., and B. Liu, "A Generic Autonomic Signaling Protocol (GRASP)", draft-ietf-anima-grasp-01 (work in progress), October 2015.

[RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, DOI 10.17487/RFC3633, December 2003, <http://www.rfc-editor.org/info/rfc3633>.

11.2. Informative References

[I-D.behringer-anima-reference-model] Behringer, M., Carpenter, B., Eckert, T., Ciavaglia, L., Liu, B., Jeff, J., and J. Strassner, "A Reference Model for Autonomic Networking", draft-behringer-anima-reference-model-04 (work in progress), October 2015.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, DOI 10.17487/RFC2629, June 1999, <http://www.rfc-editor.org/info/rfc2629>.

[RFC7575] Behringer, M., Pritikin, M., Bjarnason, S., Clemm, A., Carpenter, B., Jiang, S., and L. Ciavaglia, "Autonomic Networking: Definitions and Design Goals", RFC 7575, DOI 10.17487/RFC7575, June 2015, <http://www.rfc-editor.org/info/rfc7575>.

[RFC7576] Jiang, S., Carpenter, B., and M. Behringer, "General Gap Analysis for Autonomic Networking", RFC 7576, DOI 10.17487/RFC7576, June 2015, <http://www.rfc-editor.org/info/rfc7576>.

Authors' Addresses

Sheng Jiang (editor)
Huawei Technologies Co., Ltd
Q14, Huawei Campus, No.156 Beiqing Road
Hai-Dian District, Beijing, 100095
P.R. China

Email: jiangsheng@huawei.com

Zongpeng Du
Huawei Technologies Co., Ltd
Q14, Huawei Campus, No.156 Beiqing Road
Hai-Dian District, Beijing, 100095
P.R. China

Email: duzongpeng@huawei.com

Brian Carpenter
Department of Computer Science
University of Auckland
PB 92019
Auckland 1142
New Zealand

Email: brian.e.carpenter@gmail.com

Qiong Sun
China Telecom
No.118, Xizhimennei Street
Beijing 100035
P. R. China

Email: sunqiong@ctbri.com.cn