

ANIMA
Internet-Draft
Intended status: Standards Track
Expires: December 20, 2015

M. Behringer
Cisco Systems
June 18, 2015

An Autonomic IPv6 Addressing Scheme
draft-behringer-anima-autonomic-addressing-01

Abstract

This document describes a generic IPv6 addressing scheme which is suitable for autonomic nodes, where node addressing must not depend on a centrally managed scheme. It assumes a unique domain name and device name, and automatically derives a unique IPv6 address from those. The scheme allows for a flat address hierarchy as well as optionally, when required, the definition of zones which are aggregatable.

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 20, 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.	Requirements and Assumptions	2
3.	Fundamental Concepts of Autonomic Addressing	3
4.	The Base Addressing Scheme	4
5.	Possible Sub-Schemes	4
5.1.	Sub-Scheme 1	5
5.2.	Sub-Scheme 2	5
6.	Address Hierarchy	6
7.	IANA Considerations	7
8.	Security Considerations	7
9.	Acknowledgements	7
10.	References	7
	Author's Address	7

[1.](#) Introduction

In an Autonomic Network, as defined in [[I-D.irtf-nmrg-autonomic-network-definitions](#)], and specified in more detail in [[I-D.behringer-anima-reference-model](#)], one of the design goals is to minimise central functions. Address management traditionally is a centralised function, where addresses are assigned to nodes. In this document we assume that each node has already been assigned with a unique node name, and a domain name. We introduce an addressing scheme and an algorithm that allows the calculation of a unique IPv6 ULA address inside a domain. In other words, once a device has a unique node and domain name, this addressing scheme and algorithm allows for distributed self-management of addressing inside a network.

The addressing scheme described here is specifically designed for both the data plane of a network, as well as the Autonomic Control Plane (ACP; see [[I-D.behringer-autonomic-control-plane](#)]). It is for communication inside the domain only, specifically to support self-management functions.

[2.](#) Requirements and Assumptions

An autonomic addressing scheme has the following requirements:

- o Zero-touch for simple networks: Simple networks should have complete self-management of addressing, and not require any central address management, tools, or address planning.

Behringer

Expires December 20, 2015

[Page 2]

- o Low-touch for complex networks: If complex networks require operator input for autonomic address management, it should be limited to high level guidance only, expressed in Intent.
- o Flexibility: The addressing scheme must be flexible enough for nodes to be able to move around, for the network to grow, split and merge.
- o Robustness: It should be as hard as possible for an administrator to negatively affect addressing (and thus connectivity) in the autonomic context.
- o Support for virtualization: Autonomic Nodes may support Autonomic Service Agents in different virtual machines or containers. The addressing scheme should support this architecture.
- o Simplicity: To make engineering simpler, and to give the human administrator an easy way to trouble-shoot autonomic functions.
- o Scale: The proposed scheme should work in any network of any size.
- o Upgradability: The scheme must be able to support different addressing concepts in the future.

We assume that each node has a unique domain name and a unique node name inside that domain, for example "node17.domain.com".

3. Fundamental Concepts of Autonomic Addressing

IPv6 only: Autonomic processes should use exclusively IPv6, for simplicity reasons.

Usage: Autonomic addresses are exclusively used for self-management functions inside a trusted domain. They are not used for user traffic. Communications with entities outside the trusted domain use another address space, for example normally managed routable address space.

Separation: Autonomic address space is used separately from user address space and other address realms. This supports the robustness requirement.

Loopback-only: Only loopback interfaces of autonomic nodes carry a routable address; all other interfaces exclusively use IPv6 link local for autonomic functions. The usage of IPv6 link local addressing is discussed in [[RFC7404](#)].

Behringer

Expires December 20, 2015

[Page 3]

Use-ULA: For loopback interfaces of autonomic nodes, we use Unique Local Addresses (ULA), as specified in [RFC4193]. An alternative scheme was discussed, using assigned ULA addressing. The consensus was to use standard ULA, because it was deemed to be sufficient.

No external connectivity: They do not provide access to the Internet. If a node requires further reaching connectivity, it should use another, traditionally managed address scheme in parallel.

4. The Base Addressing Scheme

The Base ULA addressing scheme for autonomic nodes has the following format:



Figure 1: Base Addressing Scheme

The first 48 bits follow the ULA scheme, as defined in [RFC4193], to which a type field is added:

- o "FD" identifies a locally defined ULA address.
- o The "global ID" is set here to be a hash of the domain name, which results in a pseudo-random 40 bit value. It is calculated as the first 40 bits of the MD5 hash of the domain name, in the example "domain.com".
- o Type: Set to 000 (3 zero bits). This field allows different address sub-schemes in the future. The goal is to start with a minimal number of sub-scheme initially, but to allow for extensions later if and when required. This addresses the "upgradability" requirement. Assignment of types for this field should be maintained by IANA.

5. Possible Sub-Schemes

The sub-schemes listed here are not intended to be all supported initially, but are listed for discussion. The final document should define ideally only a single sub-scheme for now, and leave the other "types" for later assignment.

Behringer

Expires December 20, 2015

[Page 4]

5.1. Sub-Scheme 1




Figure 2: Addressing Scheme 1

The fields are defined as follows: [Editor's note: The lengths of the fields is for discussion.]

- o Zone ID: If set to all zero bits: Flat addressing scheme. Any other value indicates a zone. See section [Section 6](#) on how this field is used in detail.
- o Device ID: A unique value for each device.

The device ID is derived as follows: In an Autonomic Network, a registrar is enrolling new devices. As part of the enrolment process the registrar assigns a number to the device, which is unique for this registrar, but not necessarily unique in the domain. The 64 bit device ID is then composed as:

- o 48 bit: Registrar ID, a number unique inside the domain that identifies the registrar which assigned the name to the device. A MAC address of the registrar can be used for this purpose.
- o 16 bit: Device ID, a number which is unique for a given registrar, to identify the device. This can be a sequentially assigned number.

The "device ID" itself is unique in a domain (i.e., the Zone-ID is not required for uniqueness). Therefore, a device can be addressed either as part of a flat hierarchy (zone ID = 0), or with an aggregation scheme (any other zone ID). A address with zone-ID could be interpreted as an identifier, with another zone-ID as a locator.

5.2. Sub-Scheme 2




Figure 3: Addressing Scheme 2

Behringer

Expires December 20, 2015

[Page 5]

The fields are defined as follows: [Editor's note: The lengths of the fields is for discussion.]

- o Zone ID: As in sub-scheme 1.
- o Device ID: As in sub-scheme 1.
- o V: Virtualization bit(s): 1 or more bits that indicate a virtual context on an autonomic node.

In addition the scheme 1 ([Section 5.1](#)), this scheme allows the direct addressing of specific virtual containers / VMs on an autonomic node. An increasing number of hardware platforms have a distributed architecture, with a base OS for the node itself, and the support for hardware blades with potentially different OSs. The VMs on the blades could be considered as separate autonomic nodes, in which case it would make sense to be able to address them directly. Autonomic Service Agents (ASAs) could be instantiated in either the base OS, or one of the VMs on a blade. This addressing scheme allows for the easy separation of the hardware context.

The location of the V bit(s) at the end of the address allows to announce a single prefix for each autonomic node, while having separate virtual contexts addressable directly.

6. Address Hierarchy

The "zone ID" allows for the definition of a simple address hierarchy. If set to zero, the address scheme is flat. In this case, the addresses primarily act as identifiers for the nodes. Used like this, aggregation is not possible.

If aggregation is required, the 13 bit value allows for up to 8191 zones. (Theoretically, the 13 bits for the zone ID would allow also for two levels of zones, introducing a sub-hierarchy. We do not think this is required at this point, but a new type could be used in the future to support such a scheme.)

Another way to introduce hierarchy is to use sub-domains in the naming scheme. The node names "node17.subdomainA.domain.com" and "node4.subdomainB.domain.com" would automatically lead to different ULA prefixes, which can be used to introduce a routing hierarchy in the network, assuming that the subdomains are aligned with routing areas.

Behringer

Expires December 20, 2015

[Page 6]

7. IANA Considerations

The type field in the base addressing scheme should be maintained by IANA.

8. Security Considerations

tbc

9. Acknowledgements

The following people have been involved in developing this scheme: Toerless Eckert, Steinthor Bjarnason, BL Balaji, Ravi Kumar Vadapalli.

10. References

- [I-D.behringer-anima-reference-model]
Behringer, M., Carpenter, B., Eckert, T., Ciavaglia, L., and B. Liu, "A Reference Model for Autonomic Networking", [draft-behringer-anima-reference-model-02](#) (work in progress), June 2015.
- [I-D.behringer-autonomic-control-plane]
Behringer, M., Bjarnason, S., BL, B., and T. Eckert, "An Autonomic Control Plane", [draft-behringer-autonomic-control-plane-00](#) (work in progress), June 2014.
- [I-D.irtf-nmrg-autonomic-network-definitions]
Behringer, M., Pritikin, M., Bjarnason, S., Clemm, A., Carpenter, B., Jiang, S., and L. Ciavaglia, "Autonomic Networking - Definitions and Design Goals", [draft-irtf-nmrg-autonomic-network-definitions-07](#) (work in progress), March 2015.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", [RFC 4193](#), October 2005.
- [RFC7404] Behringer, M. and E. Vyncke, "Using Only Link-Local Addressing inside an IPv6 Network", [RFC 7404](#), November 2014.

Author's Address

Michael H. Behringer
Cisco Systems
Building D, 45 Allee des Ormes
Mougins 06250
France

Email: mbehring@cisco.com