Real Life Use Cases and Challenges When Implementing Link-local Addressing Only Networks as of RFC 7404

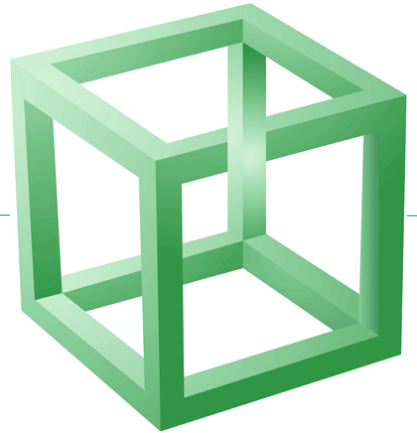

Enno Rey, erey@ernw.de

Who Am I

- Old-school network security guy with some background in provider operations.
- Involved with LIR administration in some enterprise LIRs
  - Including the one with probably the coolest org handle: ORG-HACK1-RIPE.
- IPv6 since 1999 and regularly blogging about it at `www.insinuator.net/tag/ipv6`.

ERNW
providing security.

Agenda

- Some background on RFC 7404
- Why $SOME_ORG wants to implement the approach & obstacles they've encountered
- Conclusions / Moral of the story

RFC 7404

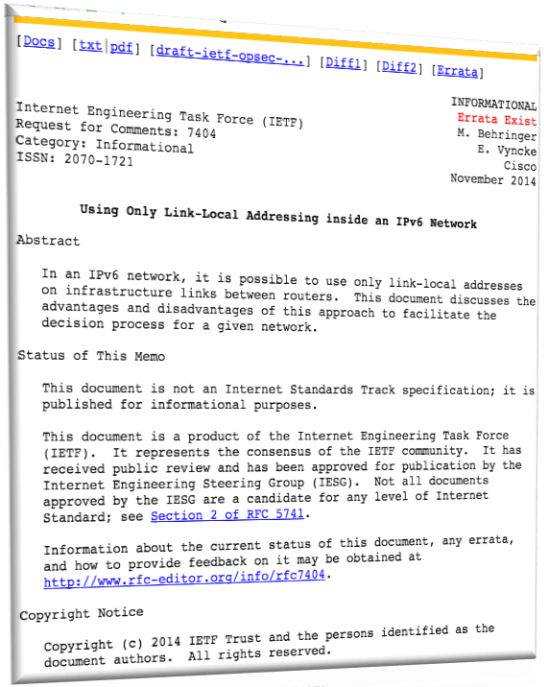

[Docs] [txt|pdf] [draft-ietf-opsec-...] [Diff1] [Diff2] [Errata]

Internet Engineering Task Force (IETF)
Request for Comments: 7404
Category: Informational
ISSN: 2070-1721

INFORMATIONAL
Errata Exist
M. Behringer
E. Vyncke
Cisco
November 2014

Using Only Link-Local Addressing inside an IPv6 Network

Abstract

In an IPv6 network, it is possible to use only link-local addresses on infrastructure links between routers. This document discusses the advantages and disadvantages of this approach to facilitate the decision process for a given network.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7404.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

- Using Only Link-Local Addressing inside an IPv6 Network
  [namely on infrastructure links]
- November 2014
- Category: Informational
  - At the time heavy discussions in OPSEC working group. RFC is supposed to discuss advantages & disadvantages, *not* to provide a recommendation.
- I for one think it's an interesting approach which can be quite beneficial in a number of use cases.

RFC 7404 – Overview of Approach

- ”Neither globally routed IPv6 addresses nor unique local addresses are configured on infrastructure links. In the absence of specific global or unique local address definitions, the default behavior of routers is to use link-local addresses, notably for routing protocols."

- Loopback interface/address assumed
  - [as source] for sending ICMPv6 messages.
  - [as destination] for management traffic.

RFC 7404 – Potential Advantages (as of RFC)

- smaller routing tables
  - and subsequently less memory consumption on routers and possibly faster convergence time
- simpler address management
- lower configuration complexity
- simpler DNS (less addresses to put into DNS)
- reduced attack surface

RFC 7404 – Potential Disadvantages (as of RFC)

- interface pings can only be performed from a node on the same link.
- traceroute (output) considered less helpful/meaningful.
- hardware dependency
- NMS tools (might need different data collection approach)

Case Study

- Very large enterprise (200K+ users, many subsidiaries) with own, wholly owned IT operations provider.
  - OEs within group = "customers".
- Company-wide MPLS network spanning several countries.
  - Main platform for PE devices is Cisco ASR 1006 & 1013 running IOS XE 03.10.
- Group level IPv6 project ongoing.

Case Study

- Identified LLA-only approach for PE-CE links, with identical addresses on all affected links, as one of the main architecture benefits of IPv6.
  - Their network, their design decisions, their (NMS) tools.
  - Trust me: they are smart people.
- Their IPAM database currently holds 43,200 networks, 20,600 (47.7%) of which are point-to-point networks.

Illustration

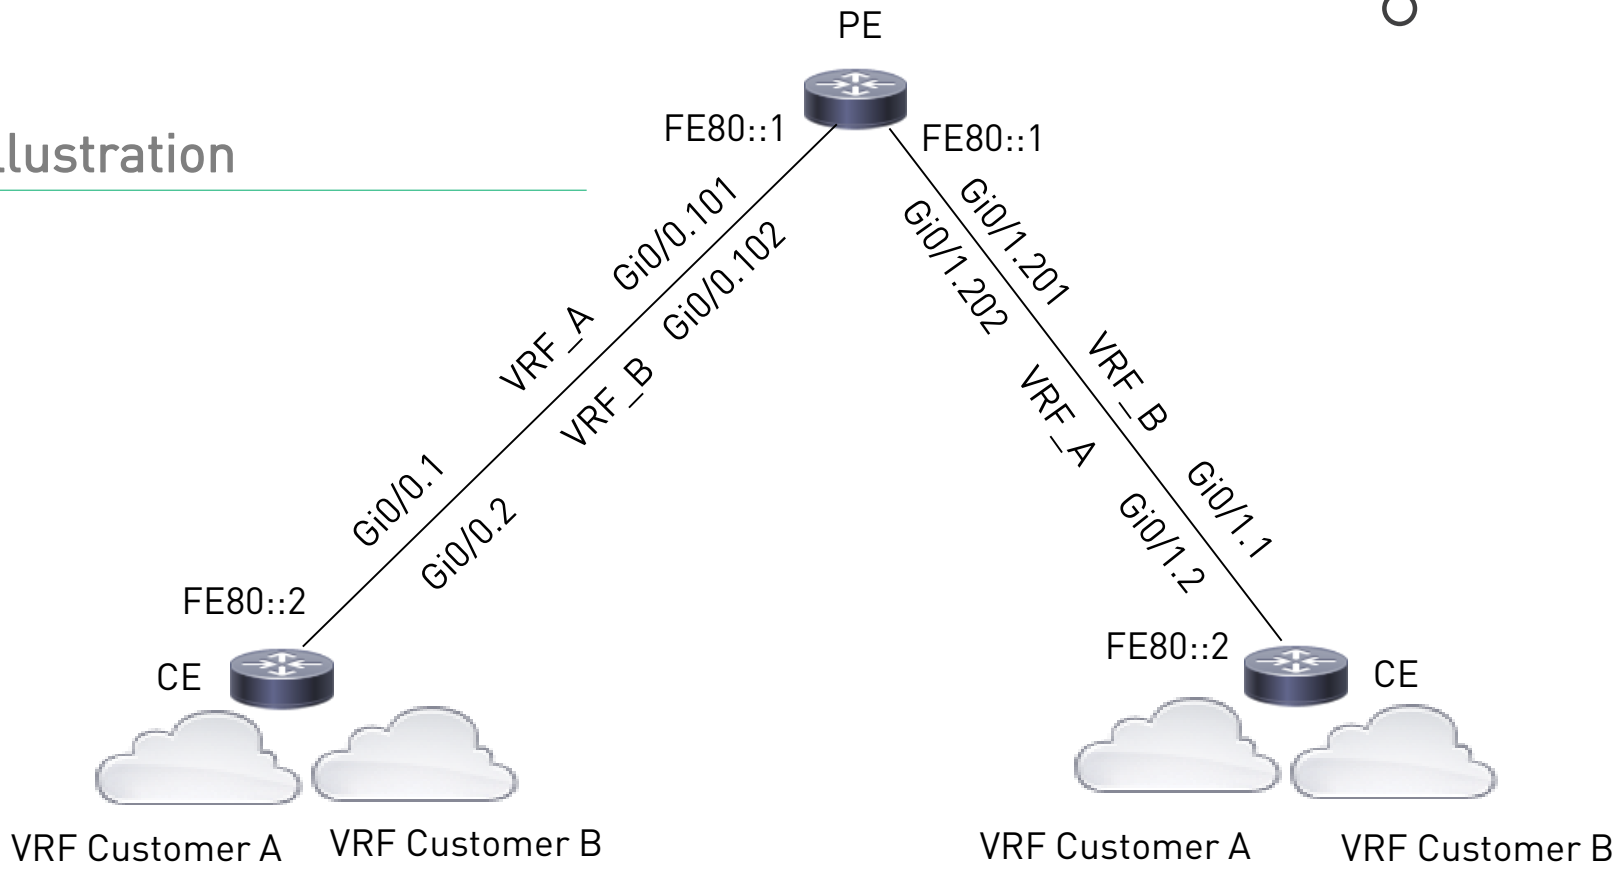

Alas...

... when performing the configuration of **second** BGP peer, the `remote-as` statement of the 1st one gets "corrupted".

So essentially the planned design & configuration approach does not work.

For reference: CSCuy05100.

```
muc-pe3(config-router-af)#neighbor FE80::2%GigabitEthernet0/0/0.4711
remote-as 65000
muc-pe3(config-router-af)#

*Jan  1 00:17:46.964: %BGP-3-NOTIFICATION: sent to neighbor
FE80::2%GigabitEthernet0/0/1 6/6 (Other Configuration Change) 0 bytes

*Jan  1 00:17:46.964: %BGP-5-NBR_RESET: Neighbor
FE80::2%GigabitEthernet0/0/1 reset (Remote AS changed)

*Jan  1 00:17:46.965: %BGP-5-ADJCHANGE: neighbor
FE80::2%GigabitEthernet0/0/1 vpn vrf customer42 Down Capability changed

*Jan  1 00:17:46.965: %BGP_SESSION-5-ADJCHANGE: neighbor
FE80::2%GigabitEthernet0/0/1 IPv6 Unicast vpn vrf customer42 topology
base removed from session  Capability changed

*Jan  1 00:17:59.391: %BGP-3-NOTIFICATION: sent to neighbor
FE80::2%GigabitEthernet0/0/1 passive 2/2 (peer in wrong AS) 2 bytes FC58

*Jan  1 00:17:59.391: %BGP-4-MSGDUMP: unsupported or mal-formatted
message received from FE80::2%GigabitEthernet0/0/1:
```

Conclusions / Moral of the Story

What does this tell us about #IPv6 in 2016?

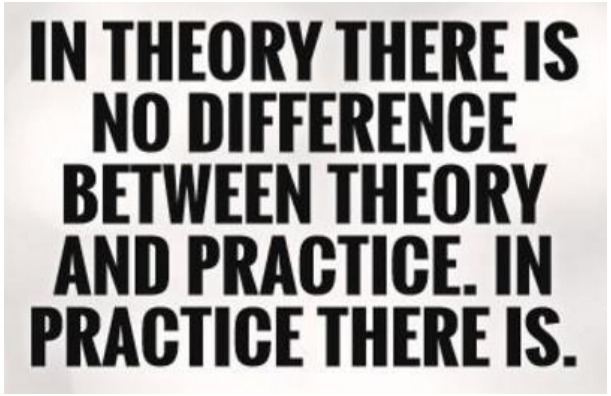

- Enterprise organizations start to realize that IPv6 can bring (not only pain & increased ops effort, but also) architecture benefits, based on paradigm shifts.
  - This is a good thing!
  - Again, I encourage you to read RFC 7404.
- There might (still) be limitations wrt vendor support though.
  - This is, well, unfortunate.
- → You *need* to test things.
  - Of course you all have well-equipped test labs, right? ;-)

ERNW
providing security.

Discussion

Do you have any questions?

There's never enough time...

THANK YOU...

...for yours!

Slides:
https://www.insinuator.net