Internet Engineering Task Force
Internet-Draft
Expires: December 28, 2006

J. Palet
C. Olvera
M. Diaz
Consulintel
June 26, 2006

# Guidelines for Numbering IPv6 Point-to-Point Links and Easing the Addressing Plans
## draft-palet-v6ops-point2point-01.txt



Abstract

This document analyzes the rational for using /64 for numbering IPv6 point-to-point links and provides some guidelines to simplify the addressing plans.

Table of Contents

## 1. Introduction

There are different alternatives for numbering IPv6 point-to-point links, and from an operational perspective, they may have different advantages or disadvantages that need to be taken in consideration under the scope of each specific network architecture design.

However, as a general rule, this document suggest the approach of using /64 in order to ensure not only compliance with standards, and consequently facilitate interoperability, but also in order to ensure avoiding possible future issues and simplifying the addressing plans.

The use of /64 also facilitates an easier way for routing the shorter aggregated prefix into the point-to-point link. Consequently it simplifies the "view" of a more unified addressing plan, providing an easier path for following up any issue when operating IPv6 networks.

The proposed approach is suitable for those point-to-point links connecting ISP to Customers, but not limited to this case, and in fact, has been tried in real scenarios for different cases. In that sense, this document can be read as guidelines for one of the possible choices available, not as a generic guideline for all the possible ways to address this.

There is another well known approach, which use two different address pools, one for the numbering the point-to-point links and another one for delegating the prefixes at the end of the point-to-point link. This document approaches for a more unified and aggregated addressing plan.

## 2. Rational for using /64

The IPv6 Addressing Architecture [1] specifies that all the Interface Identifiers for all the unicast addresses (except for 000/3) are required to be 64 bits long and to be constructed in Modified EUI-64 format. As a consequence it is forbidden to use prefixes longer than /64.

The same document also mandates the usage of the predefined subnet-router anycast address, which has cleared to zero all the bits that do not form the subnet prefix.

Moreover, [2] describes de problems of using /127 especially on point-to-point links between routers. This document also describes different choices for the point-to-point links and actually, without advocating for any specific prefix length, shows that /64 is the best solution from different perspectives, including operational

practicality.

Consequently, we shall conclude that /64 should be used for numbering point-to-point links.

## 3. Numbering Interfaces

Often, in point-to-point links, hardware tokens are not available, or there is the need to keep certain bits (u, g) cleared, so the links can be manually numbered sequentially with most of the bits cleared to zero. This numbering makes as well easier to remember the interfaces, which typically will become numbered as 1 (with 63 leading zero bits) for the provider side and 2 (with 63 leading zero bits) for the customer side.

Using interface identifies as 1 and 2 is only a very simple suggested example, and other different choices can as well be used as required in each case.

On the other hand, using the EUI-64, makes it more difficult to remember and handle the interfaces, but provides an additional degree of protection against port (actually address) scanning as described at [3].

## 4. Routing Aggregation of the Point-to-Point Links

Following this approach and assuming that a shorter prefix is typically delegated to a customer, in general a /48 [4], it is possible to simplify the routing aggregation of the point-to-point links. Towards this, the point-to-point link may be numbered using the first /64 of a given /48.

Let's see a practical example:

- A service provider uses the prefix 2001:db8::/32 and is using 2001:db8:aaaa::/48 for a given customer.

- Instead of allocating the point-to-point link from a different addressing pool, it may use 2001:db8:aaaa::/64 (which is the first /64 subnet from the 2001:db8:aaaa::/48) to number the link.

- This means that, in the case the non-EUI-64 approach is used, the point-to-point link may be numbered as 2001:db8:aaaa::1/64 for the provider side and 2001:db8:aaaa::2/64 for the customer side.

- Note that using the first /64 and interface identifiers 1 and 2 is only a very simple example, and other values may be chosen according to each case specific needs.

In this way, as the same address pool is being used for both the prefix and the point-to-point link, one of the advantages of this approach is to make very easy remembering the point-to-point links that belong to a given customer prefix, or in the other way around, remember the prefix that is linked by a given point-to-point link.

For example, making a trace-route to debug any issue to a given address in the provider network, will show a straight view, and there will not be need to check a database that related an address pool for the point-to-point links and the customer prefixes, as all they are the same.

Moreover, it is possible to use the shorter prefix as the provider side numbering for the point-to-point link and keep the /64 for the customer side. In our example, it will become:

- Point-to-point link at provider side: 2001:db8:aaaa::1/48

- Point-to-point link at customer side: 2001:db8:aaaa::2/64

This provides one additional advantage as in some platforms the configuration may be easier saving one step for the route of the delegated prefix (no need for two routes to be configured, one for the prefix, one for the point-to-point link). It is possible because the longest-prefix-match rule.

The behavior of this type of configuration has been successfully tested in different commonly available implementations with different routing protocols, including RIP, BGP, IS-IS, OSPF, along static routing, and has been used in several scenarios for a few months without any failures having been reported.

As stated in [5], "the requesting router MUST NOT assign any delegated prefixes or subnets from the delegated prefix(es) to the link through which is received the DHCP message from the delegating router", however the approach described in this document may still be useful in other DHCPv6 scenarios or non-DHCPv6 scenarios. Moreover, [6] has no explicit requirement that avoids the approach described in this document. Furthermore, this has been tested in DHCPv6-PD implementations and worked well, so we must say that it may be implementation-dependent.

## 5. Security Considerations

No security concerns seem to be related to this proposal.

## 6. IANA Considerations

This document does not have any specific IANA considerations.

## 7. Acknowledgements

The authors would like to acknowledge the inputs/comments from Alain Durand, Chip Popoviciu, Daniel Roesen, Fred Baker, Gert Doering, Olaf Bonness, Ole Troan, Pekka Savola, Vincent Jardin and the Spanish Ministry of Industry support in the co-funding of the Eureka PlaNetS project, where this work is being developed.

## 8. References

### 8.1. Normative References

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

### 8.2. Informative References

[2] Savola, P., "Use of /127 Prefix Length Between Routers Considered Harmful", RFC 3627, September 2003.

[3] Chown, T., "IPv6 Implications for TCP/UDP Port Scanning", draft-chown-v6ops-port-scanning-implications-02 (work in progress), October 2005.

[4] IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites", RFC 3177, September 2001.

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

[6] Miyakawa, S. and R. Droms, "Requirements for IPv6 Prefix Delegation", RFC 3769, June 2004.

Authors' Addresses

Jordi Palet Martinez
Consulintel
Molino de la Navata, 75
La Navata - Galapagar - Madrid
E-28420 - Spain

Phone: +34 91 151 81 99
Fax: +34 91 151 81 98
Email: jordi.palet@consulintel.es

Cesar Olvera Morales
Consulintel
Molino de la Navata, 75
La Navata - Galapagar - Madrid
E-28420 - Spain

Phone: +34 91 151 81 99
Fax: +34 91 151 81 98
Email: cesar.olvera@consulintel.es

Miguel Angel Diaz Fernandez
Consulintel
Molino de la Navata, 75
La Navata - Galapagar - Madrid
E-28420 - Spain

Phone: +34 91 151 81 99
Fax: +34 91 151 81 98
Email: miguelangel.diaz@consulintel.es

## Intellectual Property Statement



## Disclaimer of Validity



## Copyright Statement



## Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.