

# IESG Comments and Discuss Items for “draft-ietf-bmwg-ipv6-meth”

## Lars Eggert (complete)

Section 4., paragraph 2:

- > Throughout the test, the DUT can be monitored for relevant resource

Expand acronym: DUT

Section 5.3., paragraph 4:

- > extension headers (the chain MUST not contain the hop-by-hop

Nit: s/MUST not/MUST NOT/

## Amanda Baber (Complete)

Note: It would be helpful to have some language added to the IANA section or somewhere in the document about why ULA space is not appropriate and why special use space is preferred.

### Response:

This document is intended to complement and not replace RFC2544. In this sense, we tried to maintain similar ideas and principles. RFC2544 justifies the request for a distinct allocation with the argument of avoiding a collisions:

“ The network addresses 192.18.0.0 through 198.19.255.255 are have been assigned to the BMWG by the IANA for this purpose. This assignment was made to minimize the chance of conflict in case a testing device were to be accidentally connected to part of the Internet.”

### Our current text:

“ IANA reserved prefix xxxxx/48 for IPv6 benchmarking similar to 198.18.0.0/15 in RFC 3330 [9]. This prefix length provides similar flexibility as the range allocated for IPv4 benchmarking and it is taking into consideration address conservation and simplicity of usage concerns. The requested size meets the requirements for testing large network elements and large emulated networks.”

is using the same argument. RFC 2544 did not make an explicit case for not using private address space but the argument of avoiding collisions with internal, operational network traffic (of the existing organization or a future one resulting from a merger) applies in the

case of RFC 1918. The same reasoning could be used for avoiding the ULA in the case of IPv6. It is true that with ULA one could generate statistically independent address domains but that might not be perceived as sufficient. Moreover, the ULA architecture could continue to evolve as more deployment and operational experience is accumulated and that might lead to the need to update this document.

If you feel further clarification is necessary, we could insert a note in the current text such as:

“ IANA reserved prefix xxxxx/48 for IPv6 benchmarking similar to 198.18.0.0/15 in RFC 3330 [9]. This prefix length provides similar flexibility as the range allocated for IPv4 benchmarking and it is taking into consideration address conservation and simplicity of usage concerns. The requested size meets the requirements for testing large network elements and large emulated networks.

Note: Similar to RFC 2544 avoiding the use of RFC 1918 address space for benchmarking tests, this document does not recommend the use of RFC 4193 (Unique Local Addresses) in order to minimize the possibility of conflicts with operational traffic.”

#### Resolution:

Agreed to add the following note:

Note: Similar to RFC 2544 avoiding the use of RFC 1918 address space for benchmarking tests, this document does not recommend the use of RFC 4193 (Unique Local Addresses) in order to minimize the possibility of conflicts with operational traffic.

#### Jari Arkko (complete)

This is an overall good document and well worth publishing. However, there were a number of specific technical issues that should be addressed:

- > Note: During testing, either static or dynamic options for neighbor
- > discovery can be used. The static option can be used as long as it is
- > supported by the test tool. The dynamic option is preferred wherein
- > the test tool interacts with the DUT for the duration of the test to
- > maintain the respective neighbor caches in an active state.
- > To avoid neighbor solicitation (NS) and neighbor advertisement (NA)
- > storms due to the neighbor unreachability detection (NUD) mechanism
- > [3], the test scenarios assume test traffic simulates end points and
- > IPv6 source and destination addresses are one hop beyond the DUT.

Static and dynamic are not defined here or in the ND RFCs. Please define what you mean.

Response:

By static we mean a neighbor explicitly configured on the device (IPv6 address, MAC address, interface <optional or for LL>). This neighbor has an infinite lifetime. By dynamic we mean a neighbor discovered through Neighbor Discovery and one for which state is actively maintained.

Is it necessary to insert a clarification in the text?

Modified the text to include an explanation of static vs dynamic:

“Note: During testing, either static or dynamic options for neighbor discovery can be used. In the static case the IPv6 neighbor information for the test tool is manually configured on the DUT and the ipv6 neighbor information for the DUT is manually configured on the test tool. In the dynamic case, the IPv6 neighbor information is dynamically discovered by each device through the neighbor discovery protocol. The static option can be used as long as it is supported by the test tool. The dynamic option is preferred wherein the test tool interacts with the DUT for the duration of the test to maintain the respective neighbor caches in an active state. To avoid neighbor solicitation (NS) and neighbor advertisement (NA) storms due to the neighbor unreachability detection (NUD) mechanism <xref target="RFC2461"/>, the test scenarios assume test traffic simulates end points and IPv6 source and destination addresses are one hop beyond the DUT.”

Resolution: OK with the proposed text.

“Note: During testing, either static or dynamic options for neighbor discovery can be used. In the static case the IPv6 neighbor information for the test tool is manually configured on the DUT and the ipv6 neighbor information for the DUT is manually configured on the test tool. In the dynamic case, the IPv6 neighbor information is dynamically discovered by each device through the neighbor discovery protocol. The static option can be used as long as it is supported by the test tool. The dynamic option is preferred wherein the test tool interacts with the DUT for the duration of the test to maintain the respective neighbor caches in an active state. To avoid neighbor solicitation (NS) and neighbor advertisement (NA) storms due to the neighbor unreachability detection (NUD) mechanism <xref target="RFC2461"/>, the test scenarios assume test traffic simulates end points and IPv6 source and destination addresses are one hop beyond the DUT.”

Also, given the assumption about addresses on the directly attached subnets, presumably the only ND traffic, if any, would relate to address resolution of the router's address.

Response:

Our recommendation make two points:

1) Choose addresses that do not appear to be directly connected because it matches many operational cases and it avoids ND storms

2) The neighbors in this case will be at least one other device for each interface and in while both static and dynamic ND option is feasible for the small number of neighbors, we wanted to point out that test tools might not always maintain the state of the dynamically discovered neighbors and that limits the duration of tests.

**Resolution:** OK with the text change suggested bellow

> Special care needs to be taken about the neighbor unreachability  
> detection (NUD) [3] process.

What care? Be more specific. Note that as the test traffic is not from directly connected subnets, all parties involved in the test are either routers or test equipment pretending to be routers. Why would they invoke NUD?

**Response:**

I agree, this is poorly worded and we should adjust this paragraph.

Our recommendation was to statically define the neighbors during the test however, we did not identify this as a MUST. Since we do not run another control plane protocol between the DUT and the test tool (such as a routing protocol), there is no mechanism to maintain the reachability of the neighbor so NUD will be invoked. On one hand, in our tests we experienced problems with NUD properly being invoked (which led to the recommendation to use static neighbors) but on the other, should neighbors not be defined statically, we wanted to recommend consistent settings.

Most of the paragraph to which this text belongs is actually part of an earlier note:

“ Special care needs to be taken about the neighbor unreachability detection (NUD) [3] process. The IPv6 prefix reachable time in the router advertisement SHOULD be set to 30 seconds and allow 50% jitter. The IPv6 source and destination addresses SHOULD not appear to be directly connected to the DUT to avoid neighbor solicitation (NS) and neighbor advertisement (NA) storms due to multiple test traffic flows.”

So it might be better to reduce it to a note (should we think such note is useful) such as:

“When statically defined neighbors are not used, the parameters affecting Neighbor Unreachability Detection should be consistently set. The IPv6 prefix reachable time in the router advertisement SHOULD be set to 30 seconds.”

Please let us know your opinion.

**Resolution:**

Introduced the following text

“Note: When statically defined neighbors are not used, the parameters affecting Neighbor Unreachability Detection should be consistently set. The IPv6 prefix reachable time in the router advertisement SHOULD be set to 30 seconds.”

- > The IPv6 prefix reachable time in the
- > router advertisement SHOULD be set to 30 seconds and allow 50% jitter.

Jitter is not a term recognized by RFC 4861 that defines router advertisements. What specific recommendation do you have regarding the parameters that should be configured to the router advertisements?

Response:

Agreed, see above proposed text.

- > IANA reserved prefix xxxxx/48 for IPv6 benchmarking similar to
- > 198.18.0.0/15 in RFC 3330 [9]. This prefix length provides similar
- > flexibility as the range allocated for IPv4 benchmarking and it is
- > taking into consideration address conservation and simplicity of usage
- > concerns. The requested size meets the requirements for testing large
- > network elements and large emulated networks.
- >
- > Note to IANA: Replace "xxxxx" with assigned prefix.

If this is the instruction that IANA uses to make the allocation (?), it has too little instruction. Presumably you meant to say something along the lines of:

The IANA is instructed to allocate a prefix of size N from the space defined in RFC 4773 for use in ...

Response:

I believe the sole reason for choosing the current format is that the text will not require changes once the allocation is provided. We will take the approach suggested by IESG for this text.

Resolution: Modified text as described below

The IANA was instructed to allocate for IPv6 benchmarking a 48 bits prefix from the RFC 4773 pool. This allocation is similar to 198.18.0.0/15 defined in RFC 3330 [10]. This prefix length (48) provides similar flexibility as the range allocated for IPv4 benchmarking and it is taking into consideration address conservation and simplicity of usage concerns. The requested size meets the requirements for testing large network elements and large emulated networks.

Comment:

- > Tests with traffic containing each individual extension header MUST be
- > complemented with tests containing a chain of two or more extension
- > headers (the chain MUST not contain the hop-by-hop extension header).

I was surprised by the strength of the requirement here (MUST).

Overall, with the exception of hbh on routers and packet inspection features on any device, extension headers should not affect routing and forwarding performance.

Response:

When forwarding an IPv6 packet, the extension header must be traversed one by one if the upper layer information needs to be parsed as in the case of a policy defined in relation to the upper layer protocol port number. Based on field tests, while this process has an expected impact on the forwarding performance of software switched platforms, it also has a very significant impact on the forwarding performance of hardware forwarding platforms that are not designed to handle extension headers. Such platforms can regularly forward IPv6 traffic with extension headers at line rate but when an access list related to upper layer information is applied, their forwarding is not hardware assisted anymore leading to significant impact.

With our requirement, we wanted to capture this scenario which is relevant to operational environments.

Resolution: No changes made as the explanation in the text was found sufficient

- > At the end of the section in which this topic comes up, there is a
- > short explanation of these considerations. Here is the text:
- >
- > For most cases, the network elements ignore the extension headers
- > when forwarding IPv6 traffic. For these reasons it is likely the
- > performance impact related to extension headers will be observed only
- > when testing the DUT with traffic filters that contain matching
- > conditions for the upper layer protocol information. In those cases,
- > the DUT MUST traverse the chain of extension headers, a process that
- > could impact performance.
- >
- > Unless the current text is not sufficiently clear, I propose we
- > consider it to be the resolution of this point. Please let me know.

- > [permit|deny] [protocol] [SA] [DA]

The syntax for defining this is rather loose.

Response:

We tried to generalize the definition of the policy by identifying the action, the protocol and the addresses that could be part of the filter. The information types identified are detailed in the rest of that paragraph. This is similar to the filter definition in RFC 2544:

```
“ deny aa.ba.12.12 to aa.ba.112.12
  allow aa.bc.1.2 to aa.bc.65.1”
```

If this is inadequate, we could adjust the text based on your recommendations.

**Resolution:** No further changes requested.

- > where permit or deny indicates the action to allow or deny a packet
- > through the interface the filter is applied to. The protocol field is
- > defined as:
- > o ipv6: any IP Version 6 traffic
- > o tcp: Transmission Control Protocol
- > o udp: User Datagram Protocol

Does "udp" mean "udp over IPv6" in this context? Or "udp over either IPv4 or IPv6"?

**Response:**

In this document we are not targeting benchmarking of transition mechanisms such as IPv6 over IPv4 tunneling. This document provides methodology for benchmarking the forwarding of native IPv6 traffic so in this context, UDP refers to the upper layer protocol in the IPv6 packet or as you mentioned: UDP over IPv6.

**Resolution:** Agreed to add the following text at the end of the paragraph.

The filters described in these benchmarking recommendations apply to native IPv6 traffic and upper layer protocols (tcp, udp) transported in native IPv6 packets.

**Dan Romascanu (Complete)**

#### 5.1.1.1. Frame Sizes to be used on Ethernet

Ethernet in all its types has become the most commonly deployed media in today's networks. The following frame sizes SHOULD be used for benchmarking over this media type: 64, 128, 256, 512, 1024, 1280, 1518 bytes.

Note: The recommended 1518 bytes frame size represents the maximum size of an untagged Ethernet frame. According to the IEEE 802.3as standard [12], the frame size can be increased up to 2048 bytes to accommodate frame tags.

First the current increased maximal frame size accomodates not only tags but also other encapsulation required by the IEEE 802.1AE MAC security protocol. The other more significant aspect is that I believe that the recommended frame size choice should include 1522 (max size of a max length frame VLAN-tagged as per IEEE 802.1D) and 2048. My suggestion is to include these here and also in the table in Annex A.1

**Response:**

As you remember, we had a long discussion within the WG on this topic. We agreed to keep 1518 for backward compatibility and to insert the note above. The 2048 bytes long frame is included in the Annex table since it is explicitly mentioned in the note.

David Newman strongly opposed the inclusion of 1522 bytes long frame size arguing that it is outside the scope of this work to include testing with frames that have VLAN tags.

“

> I do not agree that \*any\* +4 VLAN length such as 1522 belongs in the  
> main set. Dan's comment as AD seemed to support my suggestion of  
> having a separate section for testing with VLANs.

”

While the agreement suggested by David is not evident, the discussion in the WG finally led to an agreement to drop the 1522 from the main set and to close the discussion on frame sizes with the adoption of the note above.

With this background on the WG discussions, should we still take any action on this topic?

**Resolution:**

**Updated text:**

Note: The recommended 1518 bytes frame size represents the maximum size of an untagged Ethernet frame. According to the IEEE 802.3as standard [12], the frame size can be increased up to 2048 bytes to accommodate frame tags and other encapsulation required by the IEEE 802.1AE MAC security protocol. A frame size commonly used in operational environments is 1522 bytes, max length for a VLAN-tagged frame as per 802.1D.

Added the 1522 frame size in the Table in Appendix A.1.