

Internet Engineering Task Force (IETF)
Internet-Draft
Intended status: Experimental
Expires: July 30, 2018

Khaled Omar
The Road
L. Camara
January 26, 2018

**Numbering Exchange Protocol (NEP)
Specification
draft-omar-nep-05**

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 <https://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 July 30, 2018.

Copyright Notice

Copyright (c) 2017 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 (<https://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.

Abstract

This document specifies the Numbering Exchange Protocol (NEP), a routing protocol that combines three metrics: delay, bandwidth and number of hops.

Table of Contents

1. Introduction	2
2. Numbering Exchange Protocol (NEP)	3
3. NEP Loop Prevention	5
4. Subnet Advertisement	6
5. IANA Considerations	7
6. Security Considerations	7
7. References	7
7.1. Normative References	8
7.2. Informative References	8
8. Authors' Addresses	8
Appendix A. NEP Advertisement Format	8
A.1. NEP Header	8
A.2. Topology Advertisement	9
A.3. Subnet Advertisements	10
A.3.1. IPv4 Subnet Advertisements	10
A.3.2. IPv6 Subnet Advertisements	10

1. Introduction

The Numbering Exchange Protocol (NEP) is an Interior Gateway Protocol (IGP), delivering IP packets between routers in the same Autonomous System (AS).

Many routing protocols, such as RIP, OSPF, IS-IS, and EIGRP, use normally one or even zero metrics each:

- a. RIP uses only the number of hops.
- b. OSPF uses only the cost (bandwidth).
- c. IS-IS uses only a default metric of 10 for all links.
- d. EIGRP uses the least bandwidth and delay.

NEP differs from other IGPs as it chooses its best path based on a composite metric of:

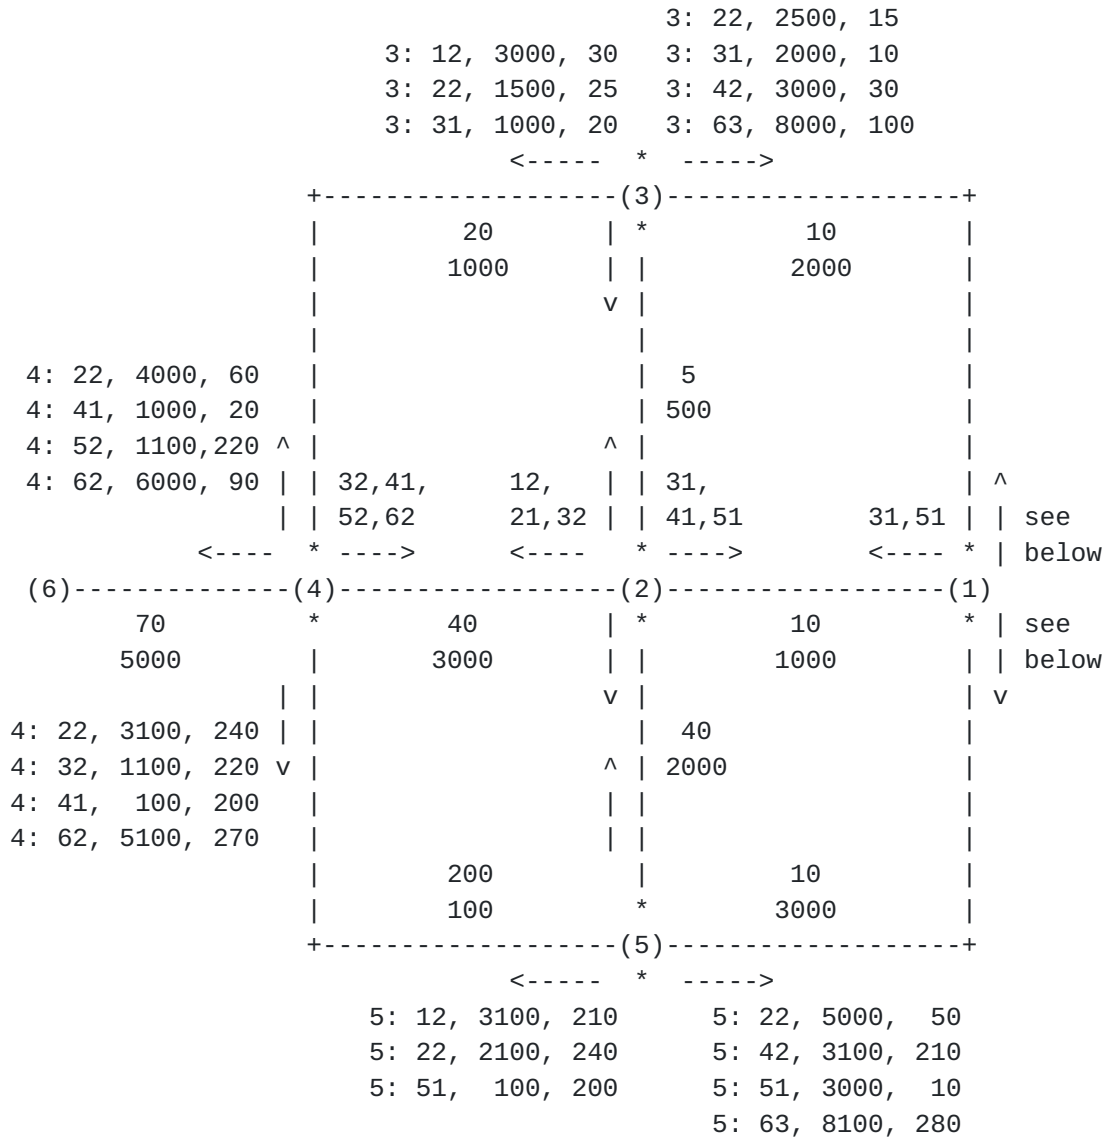
- a. Highest total bandwidth (for faster transmission).
- b. Number of hops (for less processing).
- c. Lowest total delay (for faster delivery).

NEP uses a numbering method between routers that provides a fast topology exchange and learning process, and builds a routing table shorter than those that other IGPs build.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [RFC2119, [RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

2. Numbering Exchange Protocol (NEP)

The following figure shows 6 interconnected routers within an AS as follows:



Router 1's topology advertisements cannot be represented in the figure because of the 72-column limit. Here they are:

To router 3: 1: 11, 2000, 10
 1: 22, 3000, 20
 1: 52, 5000, 20

To router 5: 1: 11, 3000, 10
 1: 22, 4000, 20
 1: 32, 5000, 20
 1: 43, 6000, 40

1: 64, 11000, 110

Omar & Camara

Expires July 30, 2018

[Page 3]

Links have two numeric values: the upper value represents the link delay, and the lower value represents the link bandwidth.

Each NEP router advertises the topology information as follows (the over-the-wire format is described in [Appendix A.1](#) and A.2):

a: rh, b, d

- a - advertising router RID.
- r - destination router RID.
- h - number of hops.
- b - total bandwidth.
- d - total delay.

Each router within an Autonomous System (AS) must be configured with a unique number called Router ID (RID).

The RID is a 32-bit decimal number that uniquely identifies a router within an AS.

Each router advertises its RID first to its neighboring routers using Hello messages.

Each router advertises all the interconnected RIDs, total bandwidth, number of hops, and total delay to the neighboring router.

Each router calculates the best path to each router using the following values:

- a. Total bandwidth.
- b. Number of hops.
- c. Total delay.

The best NEP path from the local router to every other router from the collected advertised information has the following:

- a. Highest total bandwidth.
- b. Lowest number of hops that corresponds to the highest total bandwidth.
- c. Lowest total delay.

The NEP metric is calculated using the following formula:

```
*****
*               (No. of Hops * Total-Delay)               *
*  Metric = [-----] x 10^7  *
*                   Total-BW                                *
*****
```

The best metric has the lowest value.

Each NEP router sends an Echo message of each of its interface to the neighboring router, the time taken by the message to be sent and received over the link divided by 2 determines the link delay.

The Echo message is sent by every router every 10 seconds (by default).

The link delay value is updated every 10 seconds on every advertised message that contains the 3 values that can be used to determine the best path from the NEP router to other routers.

3. NEP Loop Prevention

NEP uses a loop free mechanism by discarding the advertised topology information that has a higher metric than the existing stored entry for the same RID to prevent looping.

When an NEP router receives an advertised topology information to a specific router, it checks its own topology table, if there is no entry listed for that RID contained in the advertised message, it adds it to its own topology table.

If more than one message received to the same RID, the NEP router calculates the metric for each one of them and store the one with the lowest metric for that router.

For example:

a) Router 1 will have 3 messages to the destination router 4 as follows:

```
2: 42, 4000, 50 ==> Metric = 250,000    (discarded)
3: 42, 3000, 30 ==> Metric = 200,000    (added)
5: 42, 3100, 210 ==> Metric = 1,354,839 (discarded)
```

The message from router 3 has the lowest metric to the destination router 4 so this route through router 3 will be added to the routing table and other routes will be discarded.

b) Router 1 will have 3 messages to the destination router 6 as follows:

```
2: 63, 9000, 120 ==> Metric = 400,000    (discarded)
3: 63, 8000, 100 ==> Metric = 375,000    (added)
5: 63, 8100, 280 ==> Metric = 1,037,037 (discarded)
```

The message from router 2 has the lowest metric to the destination router 6 so this route through router 2 will be added to the routing table and other routes will be discarded.

Router 1 Routing Table:

Omar & Camara

Expires July 30, 2018

[Page 5]

```

*****
* Destination Router *      Metric      *      Next-hop Router *
*****
*          2          *      100,000    *          2          *
*****
*          3          *      50,000     *          3          *
*****
*          4          *      200,000    *          3          *
*****
*          5          *      33,333.33 *          5          *
*****
*          6          *      375,000    *          3          *
*****

```

Note: In case of two or more equal metrics for the same destination, the routes are added to the routing table and traffic is balanced between these routes.

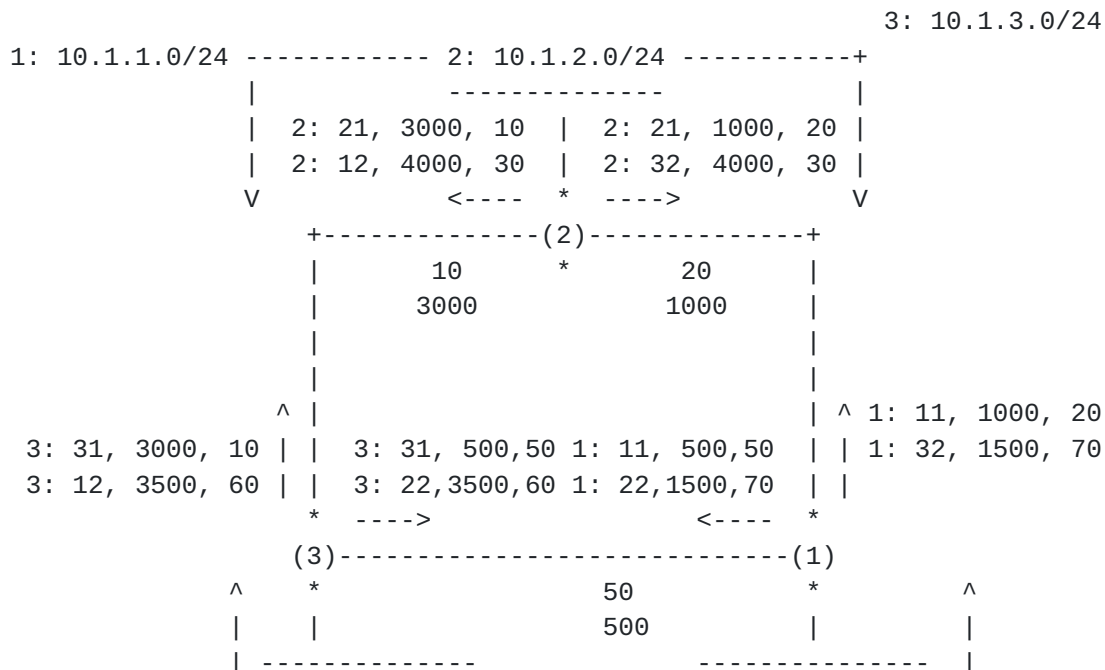
4. Subnet Advertisement

After learning the best routes from each router to other routers, every router advertises the directly connected subnet(s) to other routers.

Then, each router adds the learned subnet to the IP routing table.

Consider the following example:

Three routers are connected and every router is connected directly to one subnet.



\ 3: 10.1.3.0/24 ---> <--- 1: 10.1.1.0/24 --+
1: 10.1.1.0/24 2: 10.1.2.0/24 2: 10.1.2.0/24 3: 10.1.3.0/24

Omar & Camara

Expires July 30, 2018

[Page 6]

For router number 1, the following topology information are received:

```

2: 21, 1000, 20 ==> Metric = 200,000    (Added)
2: 32, 4000, 30 ==> Metric = 150,000    (Added)
3: 31, 500, 50  ==> Metric = 1,000,000  (Discarded)
3: 22, 3500, 60 ==> Metric = 342,857.1  (Discarded)

```

Once every router learns the best route to every other router, each router learns the advertised subnets.

The subnet information is advertised as follows (the exact over-the-wire format is specified in [Appendix A.1](#) and A.3):

```

a: subnet-number/mask  (for IPv4 [RFC791])
a: prefix/length       (for IPv6 [RFC8200])

```

Where "a" represents the advertised router RID that has the directly connected subnet or prefix.

Then, each router builds its own IP routing table:

For router number 1, the IP routing table will be as follows:

```

*****
* Dest. Router * Destination Subnet * Metric * Next-hop Router *
*****
*      2      *    10.1.2.0/24    * 200,000 *           2           *
*****
*      3      *    10.1.3.0/24    * 150,000 *           2           *
*****

```

If more than one subnet is directly connected to a router, another entry will be added to the IP routing table for that subnet with the same destination router, metric and next-hop router.

For loop prevention, each received subnet information through a specific interface is not advertised again through this interface, but out of other interfaces.

5. IANA Considerations

IANA is requested to allocate a single value from the Protocol Numbers registry, to replace TBD1 in [Appendix A.1](#).

6. Security Considerations

TBD

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), May 2017.
- [RFC8200] Deering, S., and R. Hinden, "Internet Protocol Version 6 (IPv6) Specification", STD 86, [RFC 8200](#), July 2017.

7.2. Informative References

- [RFC791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), September 1981.

[[LC to Khaled Omar: You can add your native name according to [RFC 7997](#), I will then add the substitutions. This notice is to be removed once the non-ASCII characters are added.]]

[[RFC-Editor: This document will contain non-ASCII characters when published as RFC, or when XML is used. This notice is to be removed when the non-ASCII characters are added.

s/Luis/Lu\u{00ED}s/ s/Camara/C\u{00E2}mara]]

8. Authors' Addresses

Khaled Omar Ibrahim Omar
The Road
6th of October City, Giza
Egypt

Phone: +2 01003620284
E-mail: eng.khaled.omar@hotmail.com
National ID No.: 28611262102992

Luis Camara
Portugal
E-mail: luis.camara@live.com.pt

A. NEP Advertisement Format

NEP routers advertise topology and subnet information. This appendix describes exactly how they are advertised.

Appendix A.1. NEP Header

NEP operates directly on top of IP, with protocol number TBD1.

Advertisements are communicated using the NEP header followed by the specific advertisement structure.

The structure of the NEP header is as follows:

```

          1 1 1 1 1 1 1 1 1 1 2 2 2 2 2 2 2 2 2 3 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Version # = 1 |           Type           |           Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Checksum           |           Reserved           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The version of NEP described in this document is number 1.

The Length field contains the length of the NEP data (not including the header).

The Checksum field is the one's complement of the one's complement sum of all 16-bit words of the NEP header and of the NEP data, excluding the checksum itself (or setting it to 0). If NEP is being used directly over a protocol providing other means of addressing (link/network and internetwork protocols), a pseudo-header is prepended to the NEP data.

The Reserved field is reserved for future extensibility and MUST be set to 0 by implementations not using it. [[Question: Should receivers not using the field ignore or discard packets with values of the Reserved field different from 0?]]

The pseudo-header is the concatenation of the over-the-wire representations, in the following order, of:

- o the source and destination addresses (including the lengths, if the upper protocol supports addresses of variable lengths), and
- o the protocol number for NEP in the upper protocol, padded to 32 bits, encoded as a big-endian number (in the sample captures, 143 for IP).

This algorithm is the same as used in TCP and UDP over IPv6 [[RFC8200](#)], except that it omits the length in the pseudo-header, as NEP already includes it in the NEP header.

[A.2. Topology Advertisement](#)

In Sections [2](#), [3](#) and [4](#), topology advertisements of the form "a: rh, b, d" appear. These advertisements have Type 0. Their structure is as follows:


```

      1 1 1 1 1 1 1 1 1 1 2 2 2 2 2 2 2 2 2 2 3 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Advertising Router RID                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Destination Router RID                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  # of hops  |                               Bandwidth                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Bandwidth                               |             Delay                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

[A.3. Subnet Advertisements](#)

In [Section 4](#), subnet advertisements are described. Below is their structure.

[A.3.1. Mask-Based Subnet Advertisement](#)

NEP mask-based subnet advertisements, used for IPv4 [[RFC791](#)] have the following structure and Type 1:

```

      1 1 1 1 1 1 1 1 1 1 2 2 2 2 2 2 2 2 2 2 3 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Advertising Router RID                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Subnet prefix                                       |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Subnet mask                                         |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

This structure is generic and can be adapted to NEP over any other link/network or internetwork protocol, provided that the number of octets excluding the NEP header and the advertising router RID is even: the first half of these octets are the subnet prefix and the second half are the subnet mask.

[A.3.2. Prefix-Based Subnet Advertisement](#)

NEP IPv6 [[RFC8200](#)] subnet advertisements have the following structure and Type 2:


```

          1 1 1 1 1 1 1 1 1 1 2 2 2 2 2 2 2 2 2 2 3 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|           Advertising Router RID
|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|           Prefix length
|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|
+
|
+
|           Subnet prefix
|
+
|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

The structure is generic and can be adapted to NEP over any other link/network or internetwork protocol: the subnet prefix is as many octets as an address of the link/network or internetwork protocol, and the prefix length is the number of bits in the prefix.

