           Communicating Prefix Cost to Mobile Nodes
                   draft-mccann-dmm-prefixcost-00

Abstract

   In a network implementing Distributed Mobility Management, it has
   been agreed that Mobile Nodes (MNs) should exhibit agility in their
   use of IP addresses. For example, an MN might use an old address for
   ongoing socket connections but use a new, locally assigned address
   for new socket connections. Determining when to assign a new
   address, and when to release old addresses, is currently an open
   problem. Making an optimal decision about address assignment and
   release must involve a tradeoff in the amount of signaling used to
   allocate the new addresses, the amount of utility that applications
   are deriving from the use of a previously assigned address, and the
   cost of maintaining an address that was assigned at a previous point
   of attachment. As the MN moves farther and farther from the initial
   point where an address was assigned, more and more resources are used
   to redirect packets destined for that IP address to its current
   location. The MN currently does not know the amount of resources
   used as this depends on mobility path and internal routing topology
   of the network(s) which are known only to the network operator. This
   document provides a mechanism to communicate to the MN the cost of
   maintaining a given prefix at the MN's current point of attachment so
   that the MN can make better decisions about when to release old
   addresses and assign new ones.

   This Internet-Draft will expire on September 8, 2015.

Copyright Notice

Table of Contents

[1](#). Introduction

   Previous discussions on address agility in distributed mobility
   management have focused on "coloring" prefixes with one of a small
   number of categories, such as Fixed, Sustained, or Nomadic. The
   assumption here is that the MN should use a permanent home address
   for sessions that need a persistent IP address, and a local,
   ephemeral address for short-lived sessions such as browsing.
   However, a small set of address categories lacks expressive power and
   leads to false promises being made to mobile nodes. For example, the
   concept that a home address can be maintained permanently and offered
   as an on-link prefix by any access router to which the MN may be
   attached in future is simply not attainable in the real world. There
   will always exist some access routers that do not have arrangements
   in place with the home network to re-route (via tunneling or other
   mechanisms) the home prefix to the current point of attachment.

Conversely, the assumption that a Nomadic prefix will never be
available to an MN after it changes its current point of attachment
is too limiting. There is no reason why an MN should not be able to
keep a prefix that was assigned by a first network after it moves to
a second network, provided that measures are put in place to re-route
such prefixes to the new attachment point.

Rather, this document argues that there is in reality a continuum of
cost associated with an address as the MN moves from one attachment
point to another or from one network to another. The sources of the
cost are the increased latency, network bandwidth, and network state
being maintained by a network-based mobility management scheme to
route packets destined to the prefix to the MN's current point of
attachment. By communicating this cost to the MN every time its
attachment point changes, the MN can make intelligent decisions about
when to release old addresses and when to acquire new ones.

The cost should be communicated to the MN because of several
constraints inherent in the problem:

(1) The MN is the entity that must make decisions about allocating
     new addresses and releasing old ones. This is because only the
     MN has the information about which addresses are still in use by
     applications or have been registered with other entities such as
     DNS servers.

(2) Only the network has information about the cost of maintaining
     the prefix in a network-based mobility management scheme,
     because the MN cannot know the network topology that gives rise
     to the inefficiencies.

If the cost of maintaining a prefix is not made available to the
mobile node, it may attempt to infer the cost through heuristic
mechanisms. For example, it can measure increased end-to-end latency
after a mobility event, and attribute the increased latency to a
longer end-to-end path. However, this method does not inform the MN
about the network bandwidth being expended or network state being
maintained on its behalf. Alternatively, a MN may attempt to count
mobility events or run a timer in an attempt to guess at which older
prefixes are more costly and in need of being released. However,
these methods fail because the number of mobility events is not an
indication of how far the MN has moved in a topological sense from
its original attachment point which is what gives rise to the costs
outlined above. Re-allocating an address upon expiration of a timer
may introduce uneccessary and burdensome signaling load on the
network and air interface.

1.1. Requirements Language

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 RFC 2119 [1].

1.2. Abbreviations

            MN   Mobile Node
            MPTCP Multi-Path Transmission Control Protocol
            ND   Neighbor Discovery
            PIO   Prefix Information Discovery
            SeND Secure Neighbor Discovery

2. Prefix Cost Sub-option

This document defines a prefix cost option to be carried in router
advertisements. It is a sub-option that carries meta-data as defined
by Korhonen et al. [7]

```
  0                   1                   2                   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
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |     TBD1      |       1       |C|           Reserved1         |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |          Prefix Cost          |           Reserved2          |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```
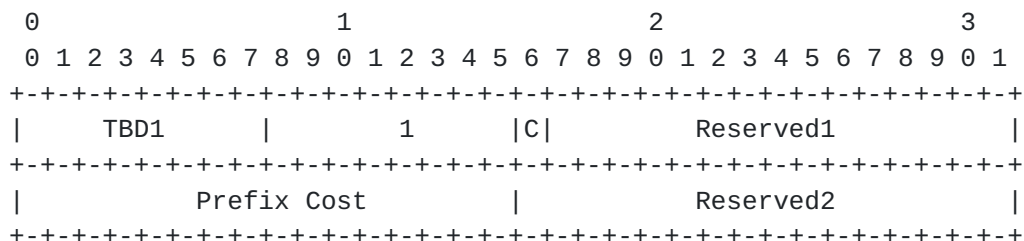
Figure 1: Prefix Cost suboption

The prefix cost is carried as a 16-bit, unsigned number in network
byte order. An higher number indicates an increased cost.

3. Host Considerations

Prefix Cost in a Router Advertisement PIO serves as a hint for the MN
to use along with application type, MN policy configuration on
network cost and available alternative routes to determine the IP
addresses and routes used. For example, if the application is
downloading a large file, it may want to maintain an IP address and
route until the download is complete. On the other hand, some
applications may use multiple connections (e.g., with MPTCP) and may
not want to maintain an IP address above a configured cost. It could
also be the case that the MN maintains the IP address even at high
cost if there is no alternative route/address. These decisions are
made based on configured policy, and interaction with applications,
all of which are internal to the MN and outside the scope of this
memo.

When the MN is ready to release an IP address, it may send a DHCPv6 [5] Release message. The network may also monitor the status of a high cost connection with Neighbor Unreachability Detection (NUD) [2], [6], and determine that an address is not used after the NUD timeout. The network should not continue to advertise this high cost route following the explicit release of the address or NUD timeout. It can initiate the release of network resources dedicated to providing the IP address to the MN.

4. Security Considerations

Security of the prefix cost option in the PIO needs to be considered. Neighbor Discovery (ND) and Prefix Information Option (PIO) security are described in [2] and [3]. A malicious node on a shared link can advertise a low cost route in the prefix cost option and cause the MN to switch. Alternatively, an incorrect higher cost route in the prefix cost option can result in the suboptimal use of network resources. In order to avoid such on-link attacks, SeND [4] can be used to reject Router Advertisements from nodes whose identities are not validated.

5. IANA Considerations

This memo defines a new Prefix Information Option (PIO) sub-option in Section 2.

6. References

6.1. Normative References

[1]   Bradner, S., "Key words for use in RFCs to Indicate
           Requirement Levels", BCP 14, RFC 2119, March 1997.

[2]   Narten, T., Nordmark, E., Simpson, W., and H. Soliman,
           "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861,
           September 2007.

[3]   Draves, R. and D. Thaler, "Default Router Preferences and
           More-Specific Routes", RFC 4191, November 2005.

[4]   Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure
           Neighbor Discovery (SEND)", RFC 3971, March 2005.

[5]   Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C.,
           and M. Carney, "Dynamic Host Configuration Protocol for
           IPv6 (DHCPv6)", RFC 3315, July 2003.

   [6]   Nordmark, E. and I. Gashinsky, "Neighbor Unreachability
                Detection Is Too Impatient", RFC 7048, January 2014.

6.2. Informative References

   [7]   Korhonen, J., Patil, B., Gundavelli, S., Seite, P., and D.
                Liu, "IPv6 Prefix Properties", draft-korhonen-6man-prefix-
                properties-02 (work in progress), July 2013.

Authors' Addresses

   Peter J. McCann (editor)
   Huawei
   400 Crossing Blvd, 2nd Floor
   Bridgewater, NJ 08807
   USA

   Phone: +1 908 541 3563
   Email: peter.mccann@huawei.com

   John Kaippallimalil
   Huawei
   5340 Legacy Dr., Suite 175
   Plano, TX 75024
   USA

   Email: john.kaippallimalil@huawei.com