Internet Engineering Task Force   Jun-ichiro itojun Hagino
INTERNET-DRAFT   IIJ Research Laboratory
Expires: January 12, 2002   Tatuya JINMEI
                                                                   Toshiba
                                                                Brian Zill
                                                                 Microsoft
                                                             July 12, 2001

Avoiding ping-pong packets on point-to-point links
draft-ietf-ipngwg-p2p-pingpong-00.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all
provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task
Force (IETF), its areas, and its working groups. Note that other groups
may also distribute working documents as Internet-Drafts.

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.''

To view the list Internet-Draft Shadow Directories, see
http://www.ietf.org/shadow.html.

Distribution of this memo is unlimited.

The internet-draft will expire in 6 months. The date of expiration will
be January 12, 2002.

Abstract

In IPv6 point-to-point link operation, there is a significant
possibility of aberrant behavior in that packets may ping-pong between
the two ends of the link. The problem can lead to wasted bandwidth and
can possibly be abused by malicious parties. This document provides an
analysis and solution to the problem.

1. The problem

In IPv6 point-to-point link operation, there are cases where we assign
an IPv6 prefix (address block) to the link. They include the following
cases:

o Link-local prefix. Each of the links has to have fe80::/64 associated
  with it. Refer to IPv6 addressing architecture [Hinden, 1998] section
  2.1 and 2.5.8.

o In some cases, we assign a global or site-local address prefix to a
  point-to-point link.

Let us assume that we have assigned P::/64 to a point-to-point link.
P::/64 could be the link-local address prefix (fe80::/64), a site-local
address prefix (fec0:0:0:x::/64) or a global address prefix. Also
assume that the Router A, which is on one end of the link, has an
address P::a. The Router B is on the other end, and has an address
P::b.

```
    Router B
      |P::b
      |
      | Point-to-point link, P::/64
      |
      |P::a
    Router A
```

In the above situation, we will see packets go ping-pong between Router
A and B under the following conditions:

o There is no link-layer address (like MAC address) for the point-to-
  point link, Link-layer address resolution procedure is not necessary.
  It is true for many of point-to-point medium, including PPP links and
  IPv6-over-IPv4 tunnels.

o A packet arrives to either of the routers, or one of the routers
  originates a packet. The destination address of the packet is P::c
  (matches P::/64, and not equal to neither P::a nor P::b).

Now, suppose that Router A is about to forward a packet, destination
address is set to P::c. The packet would leave the Router A to the
point-to-point link, as the destination address P::c satisfies the
following conditions:

o The destination address matches P::/64, and

o the destination address is not equal to P::a.

The packet reaches Router B. The address P::c is not equal to P::b, so
the Router B would forward the packet back to the point-to-point link.
The packet will be sent back to the point-to-point link repeatedly,
until the hop limit field reaches 0.

If the packet in this example had originated at Router A, it would have
a source address of P::a. This would cause Router B to hit the redirect

case given in RFC 2461, section 8.2. In this case, Router B would not
only forward the packet back to A, but it would generate a useless

redirect as well.

All IPv6 implementations are subject to the problem described in this
draft, because it is required for nodes to have a link-local address on
an interface.

1.1. Why an address prefix on a point-to-point interface?

It seems that there are a couple of different design decisions made for
existing implementations. Some of the IPv6 implementations do not
permit assigning an address prefix to a point-to-point link, and require
users to specify the whole (/128) IPv6 address on the both ends. Some
other implementations assign a prefix to a point-to-point link, and lets
the user configure an address for only one end. This is a matter of
implementation and operation choice. This document does not try to
advocate either of the implementation decisions.

Independent from the above implementation decision, there are cases
where we need to consider a point-to-point as having an IPv6 address
prefix (or prefixes).

o Link-local address (fe80::/64).

o When a node is connected to an outside network by a point-to-point
  interface, and if the node would like to use temporary addresses on
  that point-to-point link [Narten, 2001] , we must assign an IPv6
  address prefix (/64) to that link.

1.2. Why no link-layer address resolution?

When there is no link-layer address, or link-layer address resolution is
not necessary (for example, it is already configured beforehand), it is
not necessary to run Neighbor Discovery [Narten, 1998] (link-layer
address resolution). Many of the existing implementations do not. They
might run Neighbor Unreachability Detection (NUD) though, but NUD does
not really change the situation.

2. Solution

When a router attempts to forward a packet, the router SHOULD
additionally make the following check:

o Check the incoming/outgoing interface of the packet. If the interface
  is the same, is a point-to-point interface and the destination address
  on the packet seems to be on-link (in terms of Neighbor Discovery) on
  the point-to-point interface, the forwarding router SHOULD NOT forward
  the packet. Also, in this case, the router SHOULD NOT generate ICMPv6
  redirect message even if the incoming packet meets conditions in
  RFC2461 section 8.2. The router SHOULD generate an ICMPv6 error

message instead, with the type field being 1 (destination
unreachable), and the code field being 3 (address unreachable).

[3](#). Applicability statements

There is no change in node behavior, except for the error case. The
change prevents the packet from going ping-pong, and consuming the
bandwidth on the link. By implementing the change to only one of the
two routers, the ping-pong problem will go away. Therefore, the change
can be deployed gradually.

One might think that it is better to mandate NS-NA exchange (link-layer
address resolution) even for links without link-layer addresses, or
point-to-point links in general. However, from the following issues, we
have dropped the option:

o It will cause interoperability issues with existing implementations.
  If one end runs NS-NA exchange and the other end does not, they will
  not be able to exchange traffic. The authors have actually seen this
  kind of interoperability issue on an ATM PVC link during the early
  stage of IPv6 deployment.

o It would be expensive to run NS-NA exchange, just to solve this corner
  case. By adding NS-NA exchanges, it could cause packet drops during
  the "resolution" period, or at least could delay the arrival of the
  packet, while such drops and delay should not really be necessary.

[4](#). Security consideration

This draft describes a problem with point-to-point links which can be
abused by malicious parties to mount denial of service attacks from
remote locations. Such attacks could fill up point-to-links with excess
traffic. A solution to the problem is discussed and presented in the
draft.

References

Hinden, 1998.
R. Hinden and S. Deering, "IP Version 6 Addressing Architecture" in
RFC2373 (July 1998). ftp://ftp.isi.edu/in-notes/rfc2373.txt.

Narten, 2001.
T. Narten and R. Draves, "Privacy Extensions for Stateless Address
Autoconfiguration in IPv6" in RFC3041 (January 2001).
ftp://ftp.isi.edu/in-notes/rfc3041.txt.

Narten, 1998.
T. Narten, E. Nordmark, and W. Simpson, "Neighbor Discovery for IP
Version 6 (IPv6)" in RFC2461 (December 1998). ftp://ftp.isi.edu/in-
notes/rfc2461.txt.

Change history

None.

Acknowledgements

Authors' addresses

     Jun-ichiro itojun Hagino
     Research Laboratory, Internet Initiative Japan Inc.
     Takebashi Yasuda Bldg.,
     3-13 Kanda Nishiki-cho,
     Chiyoda-ku,Tokyo 101-0054, JAPAN
     Tel: +81-3-5259-6350
     Fax: +81-3-5259-6351
     Email: itojun@iijlab.net

     Tatuya JINMEI
     Research and Development Center, Toshiba Corporation
     1 Komukai Toshiba-cho, Kawasaki-shi
     Kanagawa 212-8582, JAPAN
     Tel: +81-44-549-2230
     Fax: +81-44-520-1841
     Email: jinmei@isl.rdc.toshiba.co.jp

     Brian Zill
     Microsoft Research
     One Microsoft Way
     Redmond, WA 98052
     Phone: 1-425-703-3568
     Email: bzill@microsoft.com