Using Only Link-Local Addressing inside an IPv6 Network

Abstract

   In an IPv6 network, it is possible to use only link-local addresses
   on infrastructure links between routers. This document discusses the
   advantages and disadvantages of this approach to facilitate the
   decision process for a given network.

Status of This Memo

   This document is not an Internet Standards Track specification; it is
   published for informational purposes.

   This document is a product of the Internet Engineering Task Force
   (IETF). It represents the consensus of the IETF community. It has
   received public review and has been approved for publication by the
   Internet Engineering Steering Group (IESG). Not all documents
   approved by the IESG are a candidate for any level of Internet
   Standard; see Section 2 of RFC 5741.

   Information about the current status of this document, any errata,
   and how to provide feedback on it may be obtained at
   http://www.rfc-editor.org/info/rfc7404.

Table of Contents

1. Introduction

   An infrastructure link between a set of routers typically does not
   require global or unique local addresses [RFC4193]. Using only link-
   local addressing on such links has a number of advantages; for
   example, routing tables do not need to carry link addressing and can
   therefore be significantly smaller. This helps to decrease failover
   times in certain routing convergence events. An interface of a
   router is also not reachable beyond the link boundaries, therefore
   reducing the attack surface.

   This document discusses the advantages and caveats of this approach.

   Note that some traditional techniques used to operate a network, such
   as pinging interfaces or seeing interface information in a
   traceroute, do not work with this approach. Details are discussed
   below.

   During WG and IETF last call, the technical correctness of the
   document was reviewed; however, debate exists as to whether to
   recommend this technique. The deployment of this technique is
   appropriate where it is found to be necessary.

2. Using Link-Local Addressing on Infrastructure Links

   This document discusses the approach of using only link-local
   addresses (LLAs) on all router interfaces on infrastructure links.
   Routers don't typically need to receive packets from hosts or nodes
   outside the network. For a network operator, there may be reasons to
   use addresses that are greater than link-local scope on
   infrastructure interfaces for certain operational tasks, such as
   pings to an interface or traceroutes across the network. This
   document discusses such cases and proposes alternative procedures.

2.1. The Approach

   In this approach, neither globally routed IPv6 addresses nor unique
   local addresses are configured on infrastructure links. In the
   absence of specific global or unique local address definitions, the
   default behavior of routers is to use link-local addresses, notably
   for routing protocols.

   The sending of ICMPv6 [RFC4443] error messages ("packet-too-big",
   "time-exceeded", etc.) is required for routers. Therefore, another
   interface must be configured with an IPv6 address that has a greater
   scope than link-local. This address will usually be a loopback
   interface with a global scope address belonging to the operator and
   part of an announced prefix (with a suitable prefix length) to avoid
   being dropped by other routers implementing ingress filtering
   [RFC3704]. This is implementation dependent. For the remainder of
   this document, we will refer to this interface as a "loopback
   interface".

   [RFC6724] recommends that IPv6 addresses that are greater than link-
   local scope be used as the source IPv6 address for all generated
   ICMPv6 messages sent to a non-link-local address, with the exception
   of ICMPv6 redirect messages (as defined in Section 4.5 of [RFC4861]).

   The effect on specific traffic types is as follows:

   o Most control plane protocols (such as BGP [RFC4271], IS-IS
      [IS-IS], OSPFv3 [RFC5340], Routing Information Protocol Next
      Generation (RIPng) [RFC2080], and PIM [RFC4609]) work by default
      or can be configured to work with link-local addresses.
      Exceptions are explained in the caveats section (Section 2.3).

   o Management plane traffic (such as Secure SHell (SSH) Protocol
      [RFC4251], Telnet [RFC0495], Simple Network Management Protocol
      (SNMP) [RFC1157], and ICMPv6 Echo Request [RFC4443]) can use the
      address of the router loopback interface as the destination
      address. Router management can also be done over out-of-band
      channels.

   o ICMP error messages are usually sourced from a loopback interface
      with a scope that is greater than link-local. Section 4.5 of
      [RFC4861] explains one exception: ICMP redirect messages can also
      be sourced from a link-local address.

   o Data plane traffic is forwarded independently of the link address
      type.

   o Neighbor discovery (neighbor solicitation and neighbor
      advertisement) is done by using link-local unicast and multicast
      addresses. Therefore, neighbor discovery is not affected.

   Thus, we conclude that it is possible to construct a working network
   in this way.

2.2. Advantages

   The following list of advantages is in no particular order.

   Smaller routing tables: Since the routing protocol only needs to
   carry one global address (the loopback interface) per router, it is
   smaller than the traditional approach where every infrastructure link
   address is carried in the routing protocol. This reduces memory
   consumption and increases the convergence speed in some routing
   failover cases. Because the Forwarding Information Base to be
   downloaded to line cards is smaller, and there are fewer prefixes in
   the Routing Information Base, the routing algorithm is accelerated.
   Note that smaller routing tables can also be achieved by putting
   interfaces in passive mode for the Interior Gateway Protocol (IGP).

   Simpler address management: Only loopback interface addresses need to
   be considered in an addressing plan. This also allows for easier
   renumbering.

   Lower configuration complexity: Link-local addresses require no
   specific configuration, thereby lowering the complexity and size of
   router configurations. This also reduces the likelihood of
   configuration mistakes.

   Simpler DNS: Less routable address space in use also means less
   reverse and forward mapping DNS resource records to maintain. Of
   course, if the operator selects not to enter any global interface
   addresses in the DNS anyway, then this is less of an advantage.

   Reduced attack surface: Every routable address on a router
   constitutes a potential attack point; a remote attacker can send
   traffic to that address, for example, a TCP SYN flood (see
   [RFC4987]). If a network only uses the addresses of the router
   loopback interface(s), only those addresses need to be protected from
   outside the network. This may ease protection measures, such as
   Infrastructure Access Control Lists (iACL). Without using link-local
   addresses, it is still possible to achieve the simple iACL if the
   network addressing scheme is set up such that all link and loopback
   interfaces have addresses that are greater than link-local and are
   aggregatable, and if the infrastructure access list covers that
   entire aggregated space. See also [RFC6752] for further discussion

on this topic. [RFC6860] describes another approach to hide
addressing on infrastructure links for OSPFv2 and OSPFv3 by modifying
the existing protocols. This document does not modify any protocol
and applies only to IPv6.

2.3. Caveats

   The caveats listed in this section are in no particular order.

   Interface ping: If an interface doesn't have a routable address, it
   can only be pinged from a node on the same link. Therefore, it is
   not possible to ping a specific link interface remotely. A possible
   workaround is to ping the loopback address of a router instead. In
   most cases today, it is not possible to see which link the packet was
   received on; however, [RFC5837] suggests including the interface
   identifier of the interface a packet was received on in the ICMPv6
   response. It must be noted that there are few implementations of
   this ICMPv6 extension. With this approach, it would be possible to
   ping a router on the addresses of loopback interfaces, yet see which
   interface the packet was received on. To check liveliness of a
   specific interface, it may be necessary to use other methods, such as
   connecting to the router via SSH and checking locally or using SNMP.

   Traceroute: Similar to the ping case, a reply to a traceroute packet
   would come from the address of a loopback interface, and current
   implementations do not display the specific interface the packets
   came in on. Again, [RFC5837] provides a solution. As in the ping
   case above, it is not possible to traceroute to a particular
   interface if it only has a link-local address. Conversely, this
   approach may make network topology discovery from outside the network
   simpler: instead of responding with multiple different interface IP
   addresses, which have to be correlated by the outsider, a router will
   always respond with the same loopback address. If reverse DNS
   mapping is used, the mapping is trivial in either case.

   Hardware dependency: LLAs have usually been based on 64-bit Extended
   Unique Identifiers (EUI-64); hence, they change when the Message
   Authentication Code (MAC) address is changed. This could pose a
   problem in a case where the routing neighbor must be configured
   explicitly (e.g., BGP) and a line card needs to be physically
   replaced, hence changing the EUI-64 LLA and breaking the routing
   neighborship. LLAs can be statically configured, such as fe80::1 and
   fe80::2, which can be used to configure any required static routing
   neighborship. However, this static LLA configuration may be more
   complex to operate than statically configured addresses that are
   greater than link-local scope. This is because LLAs are inherently
   ambiguous. For a multi-link node, such as a router, to deal with the

ambiguity, the link zone index must also be considered explicitly,
e.g., using the extended textual notation described in [RFC4007], as
in this example, 'BGP neighbor fe80::1%eth0 is down'.

Network Management System (NMS) toolkits: If there is any NMS tool
that makes use of an interface IP address of a router to carry out
any of its NMS functions, then it would no longer work if the
interface does not have a routable address. A possible workaround
for such tools is to use the routable address of the router loopback
interface instead. Most vendor implementations allow the
specification of loopback interface addresses for SYSLOG, IPFIX, and
SNMP. The Link Layer Discovery Protocol (LLDP) (IEEE 802.1AB-2009)
runs directly over Ethernet and does not require any IPv6 address, so
dynamic network discovery is not hindered by using only LLA when
using LLDP. But, network discovery based on Neighbor Discovery
Protocol (NDP) cache content will only display the link-local
addresses and not the addresses of the loopback interfaces;
therefore, network discovery should rather be based on the Route
Information Base to detect adjacent nodes.

MPLS and RSVP-Traffic Engineering (RSVP-TE) [RFC3209] allow the
establishment of an MPLS Label Switched Path (LSP) on a path that is
explicitly identified by a strict sequence of IP prefixes or
addresses (each pertaining to an interface or a router on the path).
This is commonly used for Fast Reroute (FRR). However, if an
interface uses only a link-local address, then such LSPs cannot be
established. At the time of writing this document, there is no
workaround for this case; therefore, where RSVP-TE is being used, the
approach described in this document does not work.

2.4. Internet Exchange Points

Internet Exchange Points (IXPs) have a special importance in the
global Internet because they connect a high number of networks in a
single location and because a significant part of Internet traffic
passes through at least one IXP. An IXP requires, therefore, a very
high level of security. The address space used on an IXP is
generally known, as it is registered in the global Internet Route
Registry, or it is easily discoverable through traceroute. The IXP
prefix is especially critical because practically all addresses on
this prefix are critical systems in the Internet.

Apart from general device security guidelines, there are basically
two additional ways to raise security (see also [BGP-OPSEC]):

1. Not to announce the prefix in question, and

2. To drop all traffic from remote locations destined to the IXP
    prefixes.

Not announcing the prefix of the IXP would frequently result in
traceroute and similar packets (required for Path MTU Discovery
(PMTUD)) being dropped due to unicast Reverse Path Forwarding (uRPF)
checks. Given that PMTUD is critical, this is generally not
acceptable. Dropping all external traffic to the IXP prefix is hard
to implement because if only one service provider connected to an IXP
does not filter correctly, then all IXP routers are reachable from at
least that service provider network.

As the prefix used in the IXP is usually longer than a /48, it is
frequently dropped by route filters on the Internet having the same
net effect as not announcing the prefix.

Using link-local addresses on the IXP may help in this scenario. In
this case, the generated ICMPv6 packets would be generated from
loopback interfaces or from any other interface with a globally
routable address without any configuration. However, in this case,
each service provider would use their own address space, making a
generic attack against all devices on the IXP harder. All of an
IXP's loopback interface addresses can be discovered by a potential
attacker with a simple traceroute; a generic attack is, therefore,
still possible, but it would require more work.

In some cases, service providers carry the IXP addresses in their IGP
for certain forms of traffic engineering across multiple exit points.
Link-local addresses cannot be used for this purpose; in this case,
the service provider would have to employ other methods of traffic
engineering.

If an Internet Exchange Point is using a global prefix registered for
this purpose, a traceroute will indicate whether the trace crosses an
IXP rather than a private interconnect. If link-local addressing is
used instead, a traceroute will not provide this distinction.

2.5. Summary

Exclusively using link-local addressing on infrastructure links has a
number of advantages and disadvantages, both of which are described
in detail in this document. A network operator can use this document
to evaluate whether or not using link-local addressing on

infrastructure links is a good idea in the context of his/her
network. This document makes no particular recommendation either in
favor or against.

3. Security Considerations

   Using only LLAs on infrastructure links reduces the attack surface of
   a router. Loopback interfaces with routed addresses are still
   reachable and must be secured, but infrastructure links can only be
   attacked from the local link. This simplifies security of control
   and management planes. The approach does not impact the security of
   the data plane. The link-local-only approach does not address
   control plane [RFC6192] attacks generated by data plane packets (such
   as hop-limit expiration or packets containing a hop-by-hop extension
   header).

   For additional security considerations, as previously stated, see
   also [RFC5837] and [BGP-OPSEC].

4. Informative References

   [BGP-OPSEC]
              Durand, J., Pepelnjak, I., and G. Doering, "BGP operations
              and security", Work in Progress, draft-ietf-opsec-bgp-
              security-05, August 2014.

   [IS-IS]   International Organization for Standardization,
              "Intermediate System to Intermediate System intra-domain
              routeing information exchange protocol for use in
              conjunction with the protocol for providing the
              connectionless-mode network service (ISO 8473)", ISO
              Standard 10589, 2002.

   [RFC0495] McKenzie, A., "Telnet Protocol specifications", RFC 495,
              May 1973, <http://www.rfc-editor.org/info/rfc0495>.

   [RFC1157] Case, J., Fedor, M., Schoffstall, M., and J. Davin,
              "Simple Network Management Protocol (SNMP)", STD 15, RFC
              1157, May 1990, <http://www.rfc-editor.org/info/rfc1157>.

   [RFC2080] Malkin, G. and R. Minnear, "RIPng for IPv6", RFC 2080,
              January 1997, <http://www.rfc-editor.org/info/rfc2080>.

   [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V.,
              and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP
              Tunnels", RFC 3209, December 2001,
              <http://www.rfc-editor.org/info/rfc3209>.

   [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed
              Networks", BCP 84, RFC 3704, March 2004,
              <http://www.rfc-editor.org/info/rfc3704>.

   [RFC4007] Deering, S., Haberman, B., Jinmei, T., Nordmark, E., and
              B. Zill, "IPv6 Scoped Address Architecture", RFC 4007,
              March 2005, <http://www.rfc-editor.org/info/rfc4007>.

   [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast
              Addresses", RFC 4193, October 2005,
              <http://www.rfc-editor.org/info/rfc4193>.

   [RFC4251] Ylonen, T. and C. Lonvick, "The Secure Shell (SSH)
              Protocol Architecture", RFC 4251, January 2006,
              <http://www.rfc-editor.org/info/rfc4251>.

   [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway
              Protocol 4 (BGP-4)", RFC 4271, January 2006,
              <http://www.rfc-editor.org/info/rfc4271>.

   [RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control
              Message Protocol (ICMPv6) for the Internet Protocol
              Version 6 (IPv6) Specification", RFC 4443, March 2006,
              <http://www.rfc-editor.org/info/rfc4443>.

   [RFC4609] Savola, P., Lehtonen, R., and D. Meyer, "Protocol
              Independent Multicast - Sparse Mode (PIM-SM) Multicast
              Routing Security Issues and Enhancements", RFC 4609,
              October 2006, <http://www.rfc-editor.org/info/rfc4609>.

   [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman,
              "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861,
              September 2007, <http://rfc-editor.org/info/rfc4861>.

   [RFC4987] Eddy, W., "TCP SYN Flooding Attacks and Common
              Mitigations", RFC 4987, August 2007,
              <http://www.rfc-editor.org/info/rfc4987>.

   [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF
              for IPv6", RFC 5340, July 2008,
              <http://www.rfc-editor.org/info/rfc5340>.

   [RFC5837] Atlas, A., Bonica, R., Pignataro, C., Shen, N., and JR.
              Rivers, "Extending ICMP for Interface and Next-Hop
              Identification", RFC 5837, April 2010,
              <http://www.rfc-editor.org/info/rfc5837>.

   [RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the
              Router Control Plane", RFC 6192, March 2011,
              <http://www.rfc-editor.org/info/rfc6192>.

   [RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown,
              "Default Address Selection for Internet Protocol Version 6
              (IPv6)", RFC 6724, September 2012,
              <http://www.rfc-editor.org/info/rfc6724>.

   [RFC6752] Kirkham, A., "Issues with Private IP Addressing in the
              Internet", RFC 6752, September 2012,
              <http://www.rfc-editor.org/info/rfc6752>.

   [RFC6860] Yang, Y., Retana, A., and A. Roy, "Hiding Transit-Only
              Networks in OSPF", RFC 6860, January 2013,
              <http://www.rfc-editor.org/info/rfc6860>.

Acknowledgments

Authors' Addresses

   Michael Behringer
   Cisco
   Building D, 45 Allee des Ormes
   Mougins 06250
   France

   EMail: mbehring@cisco.com

   Eric Vyncke
   Cisco
   De Kleetlaan, 6A
   Diegem 1831
   Belgium

   EMail: evyncke@cisco.com