

Operational Security Capabilities for
IP Network Infrastructure
Internet-Draft
Intended status: Informational
Expires: April 25, 2013

M. Behringer
E. Vyncke
Cisco
October 22, 2012

Using Only Link-Local Addressing Inside an IPv6 Network
draft-ietf-opsec-lla-only-02

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 help the decision process for a given network.

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 <http://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 April 25, 2013.

Copyright Notice

Copyright (c) 2012 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 (<http://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.

Table of Contents

1.	Introduction	3
1.1.	Requirements Language	3
2.	Using Link-Local Address on Infrastructure Links	3
2.1.	The Approach	3
2.2.	Advantages	4
2.3.	Caveats	5
2.4.	Internet Exchange Points	6
2.5.	Summary	7
3.	Security Considerations	7
4.	IANA Considerations	8
5.	Acknowledgements	8
6.	References	8
6.1.	Normative References	8
6.2.	Informative References	8
	Authors' Addresses	9

1. Introduction

An infrastructure link between a set of routers typically does not require global or even unique local addressing [[RFC4193](#)]. Using link-local addressing on such links has a number of advantages, for example that 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 horizon.

We propose to configure neither globally routable IPv6 addresses nor unique local addresses on infrastructure links of routers, wherever possible. We recommend to use exclusively link-local addresses on such links.

This document discusses the advantages and caveats of this approach.

Note: [[I-D.ietf-ospf-prefix-hiding](#)] describes another approach for OSPFv2 and OSPFv3 by modifying the existing protocols while this document does not modify any protocol but works only for IPv6.

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 [RFC2119](#) [[RFC2119](#)] when they appear in ALL CAPS. These words may also appear in this document in lower case as plain English words, absent their normative meanings.

2. Using Link-Local Address on Infrastructure Links

This document proposes to use only link-local addresses (LLA) on all router interfaces on infrastructure links. Routers typically do not need to be reached from nodes of the network, nor from outside the network. For an network operator there may be reasons to send packets to an infrastructure link for certain monitoring tasks; many of those tasks could also be handled differently, not requiring routable address space on infrastructure links.

2.1. The Approach

Neither global 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.

These link-local addresses SHOULD be hard-coded to prevent the change of EUI-64 addresses when changing of MAC address (such as after changing a network interface card).

ICMPv6 [[RFC4443](#)] error messages (packet-too-big, time-exceeded...) are required for routers, therefore a loopback interface must be configured with an IPv6 address with a greater scope than link-local (this will usually be a global scope). This greater than link-local scope IPv6 address must be used as the source IPv6 address for all generated ICMPv6 messages sent to a non link-local address and must belong to the operator and be part of an announced prefix (with a suitable prefix length) to avoid being dropped by other routers implementing [[RFC3704](#)].

The effect on specific traffic types is as follows:

- o Control plane protocols, such as BGP, ISIS, OSPFv3, RIPng, PIM work by default or can be configured to work with link-local addresses.
- o Management plane traffic, such as SSH, Telnet, SNMP, ICMP echo request ... can be addressed to loopback addresses of routers with a greater than link-local scope address. Router management can also be done over out-of-band channels.
- o ICMP error message can be sourced from a loopback address. They must not be sourced from link-local addresses when the destination is non link-local.
- 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.

We therefore conclude that it is possible to construct a working network in this way.

[2.2.](#) Advantages

Smaller routing tables: Since the routing protocol only needs to carry one loopback address per router, it is smaller than in the traditional approach where every infrastructure link addresses are carried in the routing protocol. This reduces memory consumption, and increases the convergence speed in some routing failover cases (notably because the Forwarding Information Base to be downloaded to line cards are smaller but also because there are less prefixes in

the Routing Information Base hence accelerating the routing algorithm). Note: smaller routing tables can also be achieved by putting interfaces in passive mode for the IGP.

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, or he can intent SSH brute force password attacks. If a network only uses loopback addresses for the routers, only those loopback addresses need to be protected from outside the network. This significantly eases protection measures, such as infrastructure access control lists. See also [[I-D.ietf-grow-private-ip-sp-cores](#)] for further discussion on this topic.

Lower configuration complexity: LLAs require no specific configuration (except when they are statically configured), 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 DNS mappings to maintain.

2.3. Caveats

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](#) [[RFC5837](#)] suggests to include the interface identifier of the interface a packet was received on in the ICMP response; it must be noted that there are little implementation of this ICMP extension. With this approach it would be possible to ping a router on the loopback address, yet see which interface the packet was received on. To check liveness of a specific interface it may be necessary to use other methods, for example to connect to the router via SSH and to check locally or use SNMP.

Traceroute: Similar to the ping case, a reply to a traceroute packet would come from a loopback address with a greater than link-local address. Today this does not display the specific interface the packets came in on. Also here, [RFC5837](#) [[RFC5837](#)] provides a solution.

Hardware dependency: LLAs are usually EUI-64 based, hence, they change when the MAC address is changed. This could pose 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 neighborhood. But, LLAs can be statically configured such as fe80::1 and fe80::2 which can be used to configure any required static routing neighborhood.

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

MPLS and RSVP-TE [[RFC3209](#)] allows establishing MPLS 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 Re-Route (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 proposed 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 significant part of Internet traffic pass through at least one IXP. An IXP with all the service provider nodes 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 generally two additional ways to raise security (see also [[I-D.jdurand-bgp-security](#)]):

1. Not to announce the prefix in question, and
2. To drop all traffic destined to the IXP prefixes from traffic from remote locations.

Not announcing the prefix of the IXP however would frequently result in traceroute and similar packets (required for PMTUd) to be dropped due to 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 on an IXP routes does not filter correctly, then all IXP routers are reachable from at least that service provider network.

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

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

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

2.5. Summary

Using link-local addressing only on infrastructure links has a number of advantages, such as a smaller routing table size and a reduced attack surface. It also simplifies router configurations. However, the way certain network management tasks are carried out today has to be adapted to provide the same level of detail, for example interface identifiers in traceroute.

3. Security Considerations

Using LLAs only on infrastructure links reduces the attack surface of a router: loopback addresses 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 proposal does not impact the security of the data plane. This proposal 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).

As in the traditional approach, this approach relies on the assumption that all routers can be trusted due to physical and operational security.

4. IANA Considerations

There are no IANA considerations or implications that arise from this document.

5. Acknowledgements

The authors would like to thank Salman Asadullah, Brian Carpenter, Benoit Claise, Simon Eng, Wes George, Janos Mohacsi, Alvaro Retana, Ivan Pepelnjak for their useful comments about this work.

6. References

6.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

6.2. Informative References

[I-D.ietf-grow-private-ip-sp-cores]
Kirkham, A., "Issues with Private IP Addressing in the Internet", [draft-ietf-grow-private-ip-sp-cores-07](#) (work in progress), July 2012.

[I-D.ietf-ospf-prefix-hiding]
Yang, Y., Retana, A., and A. Roy, "Hiding Transit-only Networks in OSPF", [draft-ietf-ospf-prefix-hiding-05](#) (work in progress), July 2012.

[I-D.jdurand-bgp-security]
Durand, J., Pepelnjak, I., and G. Doering, "BGP operations and security", [draft-jdurand-bgp-security-02](#) (work in progress), September 2012.

[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.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", [BCP 84](#), [RFC 3704](#), March 2004.

- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", [RFC 4193](#), October 2005.
- [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.
- [RFC5837] Atlas, A., Bonica, R., Pignataro, C., Shen, N., and JR. Rivers, "Extending ICMP for Interface and Next-Hop Identification", [RFC 5837](#), April 2010.
- [RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", [RFC 6192](#), March 2011.

Authors' Addresses

Michael Behringer
Cisco
400 Avenue Roumanille, Bat 3
Biot, 06410
France

Email: mbehring@cisco.com

Eric Vyncke
Cisco
De Kleetlaan, 6A
Diegem, 1831
Belgium

Email: evyncke@cisco.com

