v6ops Working Group	E. Levy-Abegnoli
Internet-Draft	G. Van de Velde
Intended status: Informational	C. Popoviciu
Expires: March 6, 2011	Cisco Systems
	J. Mohácsi
	NIIF/Hungarnet
	September 02, 2010

TOC

IPv6 Router Advertisement Guard

Abstract

Routed protocols are often susceptible to spoof attacks. The canonical solution for IPv6 is Secure Neighbor Discovery (SEND), a solution that is non-trivial to deploy. This document proposes a light-weight alternative and complement to SEND based on filtering in the layer-2 network fabric, using a variety of filtering criteria, including, for example, SEND status.

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 March 6, 2011.

Table of Contents

---

1. Introduction

TOC

When operating IPv6 in a shared L2 network segment without complete SEND support by all devices connected or without the availability of the infrastructure necessary to support Secure Neighbor Discovery (SEND) [RFC3971] (Arkko, J., Kempf, J., Zill, B., and P. Nikander, “SEcure Neighbor Discovery (SEND),” March 2005.), there is always the risk of facing operational problems due to rogue Router Advertisements generated maliciously or unintentionally by unauthorized or improperly configured routers connecting to the segment.
There are several examples of work done on this topic which resulted in several related studies [reference1] (LORIA/INRIA, “NDPMon - IPv6 Neighbor Discovery Protocol Monitor (http://ndpmon.sourceforge.net/),” November 2007.) [reference2] (KAME Project, “rafixd - developed at KAME

- An active rogue RA nullifier (http://www.kame.net/dev/cvsweb2.cgi/kame/kame/kame/rafixd/),” November 2007.) [reference3] (Hagino (itojun), Jun-ichiro., “Discussion of the various solutions (http://ipv6samurais.com/ipv6samurais/demystified/rogue-RA.html),” 2007.).This document describes a solution framework for the rogue-RA problem where network segments are designed around a single or a set of L2-switching devices capable of identifying invalid RAs and blocking them. The solutions developed within this framework can span the spectrum from basic (where the port of the L2 device is statically instructed to forward or not to forward RAs received from the connected device) to advanced (where a criteria is used by the L2 device to dynamically validate or invalidate a received RA, this criteria can even be based on SEND mechanisms).

---

2. Model and Applicability

TOC

RA-Guard applies to an environment where all messages between IPv6 end-devices traverse the controlled L2 networking devices. It does not apply to a shared media, when devices can communicate directly without going through an RA-Guard capable L2 networking device.
Figure 1 illustrates a deployment scenario for RA-Guard.

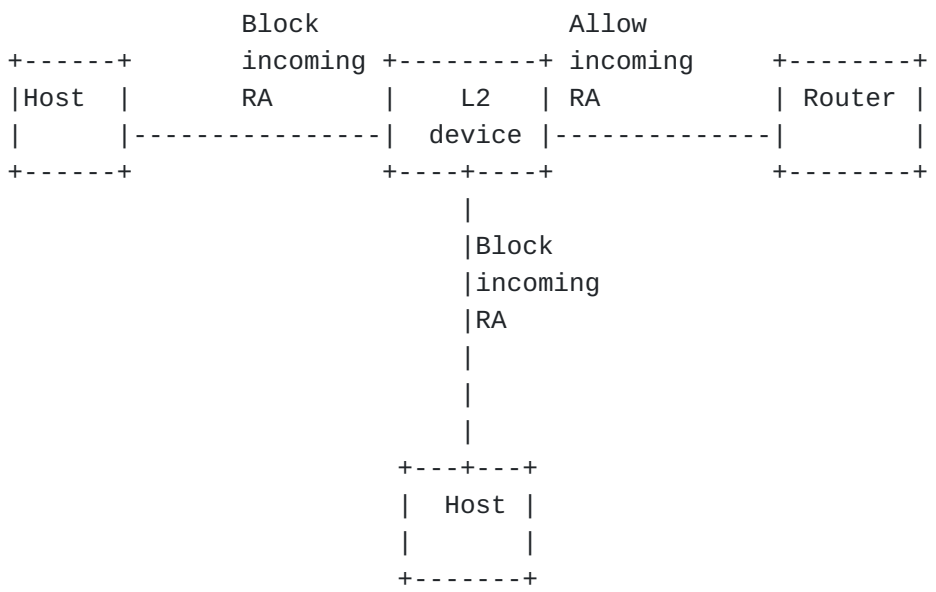

Figure 1
RA-Guard does not intend to provide a substitute for SEND based solutions. It actually intends to provide complementary solutions in those environments where SEND might not be suitable or fully supported by all devices involved. It may take time until SEND is ubiquitous in IPv6 networks and some of its large scale deployment aspects are sorted

out such as provisioning hosts with trust anchors. It is also reasonable to expect that some devices might not consider implementing SEND at all such as IPv6 enabled sensors. An RA-Guard implementation which SEND-validates RAs on behalf of hosts would potentially simplify some of these challenges.

RA-Guard can be seen as a superset of SEND with regard to router authorization. Its purpose is to filter Router Advertisements based on a set of criteria, from a simplistic "RA disallowed on a given interface" to "RA allowed from pre-defined sources" and up to full fledge SEND "RA allowed from authorized sources only".

In addition to this granularity on the criteria for filtering out Router Advertisements, RA-Guard introduces the concept of router authorization proxy. Instead of each node on the link analyzing RAs and making an individual decision, a legitimate node-in-the-middle performs the analysis on behalf of all other nodes on the link. The analysis itself is not different from what each node would do: if SEND is enabled, the RA is checked against X.509 certificates [RFC4861] (Narten, T., Nordmark, E., Simpson, W., and H. Soliman, “Neighbor Discovery for IP version 6 (IPv6),” September 2007.). If any other criteria is in use, such as known L3 (addresses) or L2 (link-layer address, port number) legitimate sources of RAs, the node-in-the middle can use this criteria and filter out any RA that does not comply. If this node-in-the-middle is a L2 device, it will not change the content of the validated RA, and avoid any of the ND-proxy pitfalls.

RA-Guard intends to provide simple solutions to the rogue-RA problem in contexts where simplicity is required while leveraging SEND in an context environment consisting of with a mix of SEND capable devices (L2 switches and routers) and devices that do not consistently use SEND. Furthermore, RA-Guard is useful to simplify SEND deployments, as only the L2 switch and the routers are required to carry certificates (their own and the trust anchor certificates).

---

3. Stateless RA-Guard

TOC

Stateless RA-Guard examines incoming RAs and decide whether to forward or block them based solely on information found in the message or in the L2-device configuration. Typical information available in the frames received, useful for RA validation is:

- *Link-layer address of the sender
- *Port on which the frame was received
- *IP source address
- *Prefix list

The following configuration information created on the L2-device can be made available to RA-Guard, to validate against the information found in the received RA frame:

*Allowed/Disallowed link-layer address of RA-sender

*Allowed/Disallowed ports for receiving RAs

*Allowed/Disallowed IP source addresses of RA-sender

*Allowed Prefix list and Prefix ranges

*Router Priority

Once the L2 device has validated the content of the RA frame against the configuration, it forwards the RA to destination, whether unicast or multicast. Otherwise, the RA is dropped.
An example of a very simple stateless RA-Guard implementation could be a small L2-switch for which there is one interface "statically-configured" as the interface connecting to a router, while all other interfaces are for non-router devices. With his small static setup the only interface forwarding RAs will be the pre-assigned router interface, while the non-router interfaces block all RAs.

---

4. Stateful RA-Guard

TOC

4.1. State Machine

TOC

Stateful RA-Guard learns dynamically about legitimate RA senders, and store this information for allowing subsequent RAs. A simple stateful scheme would be for the L2-device to listen to RAs during a certain manual determined period of time, where the start of the listening-period and the duration of the listening-period for a single instance is controled by the manual intervention. As result the L2-device can then allow subsequent RAs only on those ports on which valid RAs were received during this period. Often the LEARNING state will only be activated by manual configuration when a new IPv6 router is provisioned on the L2-network.
A more sophisticated stateful scheme is based on SEND, and is described in Section 4.2 (SEND-based RA-Guard).
The state machine for stateful RA-Guard can be global, per-interface, or per-peer, depending on the scheme used for authorizing RAs.

When RA-Guard is SEND-based, the state machine is per-peer and defined in [RFC3971].
When RA-Guard is using a discovery method, the state-machine of the RA-Guard capability consists of four different states:

*State 1: OFF

> A device or interface in RA-Guard "OFF" state, operates as if the RA-Guard capability is not available.

*State 2: LEARNING

> A device or interface in the RA-Guard "Learning" state is actively acquiring information about the IPv6 routing devices connected. The learning process takes place over a pre-defined unique period in time, set by manual configuration or it can be event triggered. The information gathered is compared against pre-defined criteria; criteria similar as the stateless RA-Guard rules to qualify the validity of the RAs.
>
> In this state, the RA-Guard enabled device or interface is either blocking "all" RAs until their validity is verified or, alternatively it can temporarily forward "all" the RAs until their validity is verified.
>
> Once the L2-device has identified through "Learning" the valid IPv6 routers and hence also identified the valid RAs, it transtions each interface from "LEARNING" into either BLOCKING state if there was no valid IPv6 router discovered at the interface, or transitions the interface into FORWARDING state if there was a valid IPv6 router discovered.

*State 3: BLOCKING

> A device or interface running RA-Guard and in Blocking state will block ingress RA-messages.
>
> An interface can transition from BLOCKING state into FORWARDING state directly if explicitly instructed by the L2-device operator.
>
> An interface can transition from BLOCKING state into LEARNING state if either explicitly told by the L2-device operator or by a triggered event.

*State 4: FORWARDING

> A device or interface running RA-Guard and in Forwarding state will accept valid ingress RAs and forward them to their destination.
>
> An interface can transition from FORWARDING state into BLOCKING state directly if explicitly instructed by the L2-device operator.
>
> An interface can transition from FORWARDING state into LEARNING state if either explicitly told by the L2-device operator or by a triggered event.

The transition between these states can be triggered by manual configuration or by meeting a pre-defined criteria.

4.2. SEND-based RA-Guard

TOC

In this scenario, the L2 device is blocking or forwarding RAs based on SEND considerations. Upon capturing an RA on the interface, the L2-device will first verify the Cryptographically Generated Addresses (CGA) [RFC3971] (Arkko, J., Kempf, J., Zill, B., and P. Nikander, “SEcure Neighbor Discovery (SEND),” March 2005.) address and the RSA (Rivest, Shamir and Adleman algorithm for public-key cryptography) signature, as specified in section 5 of [RFC3971] (Arkko, J., Kempf, J., Zill, B., and P. Nikander, “SEcure Neighbor Discovery (SEND),” March 2005.). RA should be dropped in case of failure of this verification. It will then apply host behavior as described in section 6.4.6 of [RFC3971] (Arkko, J., Kempf, J., Zill, B., and P. Nikander, “SEcure Neighbor Discovery (SEND),” March 2005.). In particular, the L2 device will attempt to retrieve a valid certificate from its cache for the public key referred to in the RA. If such certificate is found, the L2 device will forward the RA to destination. If not, the L2 device will generate a Certification Path Solicitation (CPS) [RFC3971] (Arkko, J., Kempf, J., Zill, B., and P. Nikander, “SEcure Neighbor Discovery (SEND),” March 2005.), sourced with UNSPECIFIED address, to query the router certificate(s). It will then capture the Certification Path Advertisements (CPA) [RFC3971] (Arkko, J., Kempf, J., Zill, B., and P. Nikander, “SEcure Neighbor Discovery (SEND),” March 2005.), and attempt to validate the certificate chain. Failure to validate the chain will result in dropping the RA. Upon validation success, the L2 device will forward the RA to destination and and store the router certificate in its cache.
In order to operate in this scenario, the L2-device should be provisioned with a trust anchor certificate, as specified in section 6 of [RFC3971]. It may also establish a layer-3 connectivity with a Certificate Revocation List (CRL) Certification Path Advertisement server and/or with and NTP server. Bootstrapping issue in this case can

be resolved by using the configuration method to specify a trusted port to a first router, and SEND-based RA-Guard method on all other ports. The first router can then be used for Network Time Protocol (NTP) [RFC5905] (Mills, D., Martin, J., Burbank, J., and W. Kasch, “Network Time Protocol Version 4: Protocol and Algorithms Specification,” June 2010.) and CRL connectivity.

---

5. RA-Guard Use Considerations

TOC

The RA-Guard mechanism is effective only when all messages between IPv6 devices in the target environment traverse controlled L2 networking devices. In the case of environments such as Ethernet hubs, devices can communicate directly without going through an RA-Guard capable L2 networking device, the RA-Guard feature cannot protect against rogue-RAs.
RA-Guard mechanisms do not offer protection in environments where IPv6 traffic is tunneled.

---

6. IANA Considerations

TOC

There are no extra IANA consideration for this document.

---

7. Security Considerations

TOC

Once RA-Guard has setup the proper criteria, for example, it identified that a port is allowed to receive RAs, or it identified legitimate sources of RA, or certificate base [RFC4861] (Narten, T., Nordmark, E., Simpson, W., and H. Soliman, “Neighbor Discovery for IP version 6 (IPv6),” September 2007.), then there is no possible instances of accidently filtered legitimate Router advertisements assuming the RA-Guard filter enforcement follows strictly the RA-Guard set criteria.
in Section 4.1 a simple mechanism to learn dynamical the valid IPv6 routers connected to a L2-device is explained. It is important that this LEARN state is only entered intentionally by manual configuration. The list of learned IPv6 routers should be verified by the network manager to make sure that it corresponds with the expected valid RA list. This procedure will make sure that either accidently or intentionally rogue RAs are blocked by RA-guard.

---

8. Acknowledgements

TOC

The authors dedicate this document to the memory of Jun-ichiro Hagino (itojun) for his contributions to the development and deployment of IPv6.

---

9. References

TOC

9.1. Normative References

TOC

[RFC3971]	Arkko, J., Kempf, J., Zill, B., and P. Nikander, “SEcure Neighbor Discovery (SEND),” RFC 3971, March 2005 (TXT).
[RFC4861]	Narten, T., Nordmark, E., Simpson, W., and H. Soliman, “Neighbor Discovery for IP version 6 (IPv6),” RFC 4861, September 2007 (TXT).
[RFC4158]	Cooper, M., Dzambasow, Y., Hesse, P., Joseph, S., and R. Nicholas, “Internet X.509 Public Key Infrastructure: Certification Path Building,” RFC 4158, September 2005 (TXT).
[RFC5905]	Mills, D., Martin, J., Burbank, J., and W. Kasch, “Network Time Protocol Version 4: Protocol and Algorithms Specification,” RFC 5905, June 2010 (TXT).

9.2. Informative References

TOC

[reference1]	LORIA/INRIA, “NDPMon - IPv6 Neighbor Discovery Protocol Monitor (http://ndpmon.sourceforge.net/),” November 2007.
[reference2]	KAME Project, “rafixd - developed at KAME - An active rogue RA nullifier (http://www.kame.net/dev/cvsweb2.cgi/kame/kame/kame/rafixd/),” November 2007.
[reference3]	Hagino (itojun), Jun-ichiro., “Discussion of the various solutions (http://ipv6samurais.com/ipv6samurais/demystified/rogue-RA.html),” 2007.
[reference4]	Chown, Tim. and Stig. Venaas, “Rogue IPv6 Router Advertisement Problem (draft-ietf-v6ops-rogue-ra-00.txt),” May 2009.

Authors' Addresses

TOC

	Eric Levy Abegnoli
	Cisco Systems
	Village d'Entreprises Green Side - 400, Avenue Roumanille
	Biot - Sophia Antipolis, PROVENCE-ALPES-COTE D'AZUR 06410
	France
Phone:	+33 49 723 2620
Email:	elevyabe@cisco.com
	Gunter Van de Velde
	Cisco Systems
	De Kleetlaan 6a
	Diegem 1831
	Belgium
Phone:	+32 2704 5473
Email:	gunter@cisco.com
	Ciprian Popoviciu
	Cisco Systems
	7025-6 Kit Creek Road
	Research Triangle Park, North Carolina NC 27709-4987
	USA
Phone:	+1 919 392-3723
Email:	cpopovic@cisco.com
	János Mohácsi
	NIIF/Hungarnet
	18-22 Victor Hugo
	Budapest H-1132
	Hungary
Phone:	tbc
Email:	mohacsi@niif.hu