| | | |
|---|---|---|
| v6ops Working Group | E. Levy-Abegnoli | |
| Internet-Draft | G. Van de Velde | |
| Expires: March 14, 2009 | C. Popoviciu | |
| | Cisco Systems | |
| | J. Mohácsi | |
| | NIIF/Hungarnet | |
| | September 10, 2008 | |

TOC

**IPv6 RA-Guard**

**Status of this Memo**



**Abstract**

It is particularly easy to experience "rogue" routers on an unsecured link. Devices acting as a rougue router may send illegitimate RAs. Section 6 of SeND [RFC3971] (Arkko, J., Kempf, J., Zill, B., and P. Nikander, “SEcure Neighbor Discovery (SEND),” March 2005.) provides a full solution to this problem, by enabling routers certification. This solution does, however, require all nodes on an L2 network segment to support SeND, as well as it carries some deployment challenges. End-nodes must be provisioned with certificate anchors. The solution works better when end-nodes have access to a Certificate Revocation List server, and to a Network Time Protocol server, both typically off-link, which brings some bootstrap issues.
When using IPv6 within a single L2 network segment it is possible and sometimes desirable to enable layer 2 devices to drop rogue RAs before

they reach end-nodes. In order to distinguish valid from rogue RAs, the L2 devices can use a spectrum of criterias, from a static scheme that blocks RAs received on un-trusted ports, or from un-trusted sources, to a more dynamic scheme that uses SeND to challenge RA sources.
This document reviews various techniques applicable on the L2 devices to reduce the threat of rogue RAs.

---

**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 SeND, 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 such as an Ethernet hub, when devices can communicate directly without going through an RA-Guard capable L2 networking device.
Figure 1 illustrates a deployment scenario for RA-Guard.

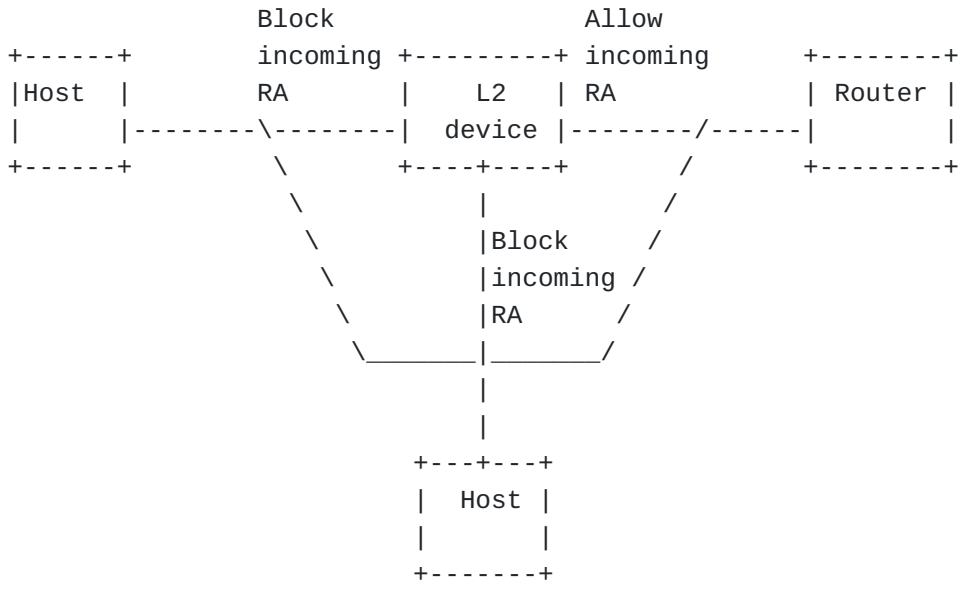


RA-Guard does not intend to provide a substitute for SeND based solutions. It actually intends to provide complementary solutions in those environments where SeND might not be suitable or fully supported by all devices involved. It may take time until SeND is ubiquitous in IPv6 networks and some of its 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 criterias, from a simplistic "RA disallowed on a given

interface" to "RA allowed from pre-defined sources" and up to full fladge 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. 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.

---

## 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 period of time, then to allow subsequent RAs only on those ports on which valid RAs were received during this period. 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 devices connected to its interfaces. The learning process takes place over a pre-defined period of time by capturing router advertisements or it can be event triggered. The information gathered is compared against pre-defined criteria which 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 the RAs until the decision is being made.

*State 3: BLOCKING

> A device or interface running RA-Guard and in Blocking state will block ingress RA-messages.

*State 4: FORWARDING

> A device or interface running RA-Guard and in Forwarding state will accept ingress RAs and forward them to their destination/

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 CGA address and the RSA signature, as specified in section 5 of [RFC3971]. 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]. 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 CPS, sourced with UNSPECIFIED address, to query the router certificate(s). It will then capture the CPA(s), 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 CRL 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 NTP and CRL connectivity.

---

TOC

## 5. RA-Guard Use Considerations

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

There are no extra Security consideration for this document.

---

## 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. 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). |
| [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. |


## 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 |


## Full Copyright Statement

TOC

**Intellectual Property**