

Home Networking
Internet-Draft
Intended status: Informational
Expires: May 10, 2013

E. Kline
Google Japan
November 6, 2012

Default Border Definition
draft-kline-default-perimeter-01

Abstract

Automatic, simple identification of when traffic is crossing a perimeter is highly desirable for a variety of home network uses. This document describes how to use homenet routing protocol adjacencies as the primary signal of a common administrative domain (e.g. "the home"). Classification of interfaces et cetera as internal or external follow from this, as do various policy and implementation implications. One fundamental implication is that the active definition of a home network's interior is no more secure than its policy for forming homenet routing protocol adjacencies.

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 May 10, 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. Terminology	3
3. Dynamically determining the border	4
3.1. Collect information about neighboring routers	5
3.2. Classify next hops	5
3.3. Classify interfaces	6
3.3.1. Mixed mode	6
3.4. State changes and logging	6
4. Fixed-category interfaces	6
4.1. Examples	7
5. Security Considerations	7
5.1. Disclamatory remarks	7
5.2. Security of adjacency formation	8
5.2.1. Secure links and authenticated adjacency formation	8
5.2.2. Unsecure links	8
5.2.3. Recommendations	9
5.3. Example border-aware filtering policies	9
5.3.1. Anti-spoofing on internal interfaces	9
5.3.2. Stateful filtering on external interfaces	10
5.3.3. Mixed-mode interface filtering	10
6. Acknowledgements	11
7. IANA Considerations	11
8. References	11
8.1. Normative References	11
8.2. Informative References	11
Author's Address	12

1. Introduction

Automatic, simple identification of when traffic is crossing the homenet perimeter is highly desirable for a variety of home network uses. This is a non-trivial task as it is tantamount to automated discovery of administrative boundaries.

Many architectures make it difficult or impossible (by design) to detect administrative boundaries and rely on various mechanisms of user or administrator invention to create or modify such boundaries. Other hints about administrative boundaries can be insecure, unreliable, operationally impractical, or may place arbitrary requirements upon the architecture where previously no such requirement existed.

This document describes how to use homenet routing protocol adjacencies as the primary signal of a common administrative domain (e.g. "the home"). Classification of interfaces et cetera as internal or external follow from this, as do various policy and implementation implications. One fundamental implication is that the active definition of a home network's interior is no more secure than its policy for forming homenet routing protocol adjacencies.

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 RFC 2119 [RFC2119].

2. Terminology

In order to address border determination at a manageable scale the scope has been limited to discussing concepts of "interior", "exterior", and "border". Documents may reference any of the terms "internal", "interior", "external", or "exterior" as required by grammar (adjective versus noun use cases). Definitions in use by this document are as follows.

Internal/Interior

The interior is broadly defined to include the collection of interfaces (physical or virtual), nodes, and forwarding next hops collectively under the control of a single logical administrative domain. This document uses the homenet routing protocol adjacencies as a indicator of membership in the same logical administrative domain.

External/Exterior

The exterior is broadly defined to include all interfaces (physical or virtual), nodes, and forwarding next hops collectively NOT under the control of any single logical administrative domain and specifically NOT under the control of the administrative domain which defines the interior.

Border/Perimeter

The border is taken to be the collection of all ephemeral demarcations within the collection of interior nodes which forward traffic such that any IP packet transiting that demarcation can be said to be crossing either from the interior toward the exterior or from the exterior toward the interior. This is independent of whether or not such traffic is permitted by policy to complete its transiting from one zone to the other.

Additionally, some implementations MAY choose to support handling questionable home network configurations which result in an interface qualifying for both interior and exterior classification simultaneously. Requirements for this are discussed further below.

Expressly not considered by this document are architectures having multiple interior networks, nor the relationships between them as separate from their relationship to their common exterior or any common border (e.g. a hierarchy of internal networks). This document is solely concerned with a single interior, a single exterior, and a single logical perimeter between the two.

3. Dynamically determining the border

Homenet routers may support interfaces which attempt to learn the nature of their relationship to neighboring routers, and determine where the border between the "interior" and the "exterior" lies.

Interfaces which have not yet determined their categorization may consider themselves to be in a "learning" state, where no traffic is routed but probing continues. Once nodes of any kind (e.g. either routers or hosts) are detected, classification takes place and the router exits the "learning" state.

The classification algorithm documented here is roughly:

1. Continuously collect information on all interfaces about neighboring routers (including manually configured routers).

2. Classify next hops as either "internal" or "external" primarily by homenet router protocol adjacency status.
3. Classify interfaces according to the nature of their next hops.

Manually configured next hops MUST also have their classification as either "internal" or "external" explicitly specified. As such, they can be considered to be of a fixed category, and on-going evaluation is NOT required.

Policies of various sorts can be applied and updated as appropriate based on these classifications, as and when they are determined.

3.1. Collect information about neighboring routers

An interface (physical or virtual) which is configured to dynamically assess its internal or external classification MUST periodically probe for routing information on the link. This includes sending Router Solicitations, DHCPv6 Prefix Delegation requests, and probing for homenet routing protocol-capable nodes.

Probing for routing information MUST be performed whenever the interface is logically up. The periodicity of probes is protocol-dependent (e.g. subject to Router Advertisement lifetimes, DHCPv6 lease timers, or homenet routing protocol timers). Wherever possible, implementations SHOULD limit the impact of probing by implementing mechanisms like exponential back-off.

Homenet routers MUST be able to identify loops, e.g. when two (or more) of the router's own interfaces are connected on the same link. Compliant routers MUST administratively log this misconfiguration, and SHOULD implement a mechanism that permits maximum continued homenet functionality if possible. For example, implementations MAY administratively disable all but one of looped interfaces.

3.2. Classify next hops

Routing information is used to categorize next hops as either "internal" or "external".

Routers with which a homenet routing protocol adjacency is successfully established MUST be considered "internal".

Routers of which this homenet router has knowledge but with which no homenet routing protocol adjacency is successfully establish AND from which no routing information is learned SHOULD be considered "internal". This includes "downstream" routers for which the homenet router is acting as the Delegating Router via a DHCPv6 Prefix

Delegation exchange.

All other routers with which no homenet routing protocol adjacency is successfully established MUST be considered "external".

3.3. Classify interfaces

An interface with no "external" next hops SHOULD be categorized as "internal". This includes interfaces serving leaf networks consisting only of hosts, an interface which has "downstream" routers for which this router is a Delegating Router, an interface with only homenet routing protocol adjacent peers, or any combination thereof.

An interface with next hops all of which are categorized as "external" MUST be categorized as "external".

3.3.1. Mixed mode

Some devices MAY choose to support handling questionable home network configurations which result in an interface having both interior and exterior next hops simultaneously. This can happen if, for example, two homenet routers form an adjacency with each other over the same interface they use for communicating to "upstream" ISP routers.

All homenet routers, whether this configuration is considered supported or not, MUST administratively log and provide product-relevant notification of this configuration, preferably with recommendations for resolution.

3.4. State changes and logging

Home routers performing dynamic border discovery MUST continuously evaluate the interior and exterior classifications of interfaces. These may change at any time, for example if new devices are added into the network or a power event causes all equipment to reset, and specific ordering of device bring-up within a homenet network MAY NOT be assumed.

Homenet routers performing dynamic border discovery SHOULD administratively log the perimeter classification of all interfaces (physical and virtual), the reason(s) for such classification, and times at which such classifications are made or changed.

4. Fixed-category interfaces

Interfaces (physical or virtual) which have product-defined purposes or are otherwise permanently categorized by the homenet router

implementation as exclusively "internal" or exclusively "external" do not require any algorithm to determine their categorization.

Homenet routers MUST restrict relevant traffic on fixed-category interfaces according to their categorizations. Specifically, they MUST NOT originate traffic which could result in re-categorizing the interface IF the interface were dynamically attempting to learn its categorization. For example, a fixed "external" interface MUST NOT attempt to participate in the homenet routing protocol. Similarly, fixed "internal" interfaces must not issue DHCPv6 Prefix Delegation requests.

4.1. Examples

Examples of product-defined interfaces include home router interfaces which are labeled for their intended use, e.g. RJ-45 ports labeled "WAN" and "LAN" or symbols indicating "The Internet" and "inside the home". Other examples include wireless routers with defined separate WLAN "home" and "guest" ESSIDs.

Another set of examples of product-defined, fixed category interfaces are those which require subscriber credentials in order for that interface to successfully pass general purpose traffic. These include authenticated PPPoE sessions and 3G/LTE PDP contexts (e.g. requiring a SIM card associated with a valid customer account). These SHOULD be classified as "exterior".

Similarly, an implementation may consider the interface a mobile device uses to provide service to "tethering" clients to be a fixed-category interface. Such interfaces SHOULD be classified as "interior".

5. Security Considerations

A key motivation for border identification is the application of security policies which can take into account classifications of interior, exterior, and the transition from one the other. General remarks, comments on the security of the adjacencies which form the basis of border identification, and examples of potential policies which might be applied follow.

5.1. Disclamatory remarks

By default all network nodes SHOULD follow best current security practices. Any node may at times find itself in a hostile environment. This is obviously true of mobile nodes, for example, when connecting to any public "Wi-Fi" network. This is, of course,

equally true of more traditionally "fixed" nodes. Any compromised neighbor node ("fixed" or mobile) may be used as a conduit for further compromise. Indeed, one study of a captured "botnet" ([TORPIG], section 5.2.4) found that roughly 78.9% of infected hosts had RFC 1918 [RFC1918] addresses, commonly used in IPv4 NAT deployments.

Though it goes without saying, at all times homenet implementers MUST remain mindful of best current security practices, including but not limited to RFC 4864 [RFC4864], RFC 4890 [RFC4890], RFC 6092 [RFC6092], and others.

5.2. Security of adjacency formation

The security of the border definition is limited by the security applied to the formation of homenet routing protocol adjacencies: the next hops with which a homenet router forms adjacencies are the next hops the router trusts with the application of interior policies.

5.2.1. Secure links and authenticated adjacency formation

The trustworthiness of next hops during adjacency formation can be improved if the security of the link connecting them can be trusted. Using encrypted link technologies like 802.1x or secured "Wi-Fi" ESSIDs when forming homenet adjacencies, or authenticating homenet next hops by physical or cryptographic mechanisms limits the ability of malicious nodes to join the homenet interior.

5.2.2. Unsecure links

In general IP over Ethernet connections, common to residential Internet (and countless other places like some in-room hotel network) service provider deployments, create the possibility of malicious nodes attempting to join the homenet interior.

In a broad variety of circumstances users already implicitly trust unsecured links. Residential subscribers generally trust that their ISP has properly isolated their connection from any neighbors; few if any subscribers validate the ISP's DHCP server in order to thwart a malicious neighbor intervening.

In the event of a network with a single upstream where an interior next hop is formed instead of an external next hop, the homenet network as a whole would have detected no external next hops. Homenet router networks in which there are no external next hops SHOULD administratively log this configuration and SHOULD provide a means to alert the user to this condition. Note that for isolated networks of homenet routers (e.g. a lab network) this configuration

is entirely valid.

User notification alone is not sufficient protection for the homenet user, and will not correctly alert the user of a homenet with two upstream connections, one of which has mistakenly not categorized a next hop as external. To better serve the homenet user, homenet routers are **SHOULD** follow one or more of the recommendations in Section 5.2.3.

5.2.3. Recommendations

Homenet router implementations that support dynamic discovery of the border (i.e. have interfaces on which the dynamic border detection described in Section 3 can be configured to operate) **SHOULD** support restricting homenet routing protocol adjacency formation to only next hops which meet some user-defined or user-verified authentication mechanism (including examples described in Section 5.2.1).

Alternatively, implementations **MAY** incorporate a mechanism (possibly physical) whereby a homenet user can disable border detection on an interface which the user wishes to force into either an interior or exterior classification (e.g. a button to force an interface to be "external" only).

5.3. Example border-aware filtering policies

As a homenet router forms adjacencies and learns internal aggregate prefixes it could dynamically maintain a single logical entity encompassing all current internal prefixes in use that can be treated as a whole (e.g. an access list). Below are example filtering policies that might be applied by homenet routers with knowledge of both this prefix set and the interior or exterior classification of all interfaces.

The examples below use the string "{interior_nets}" for refer to the grouping of all internal aggregate prefixes. The sample filtering policy rules are written in configuration pseudo-syntax that should hopefully be intuitive.

5.3.1. Anti-spoofing on internal interfaces

Given knowledge of all interior network prefixes and the categorization of interfaces, all interior interfaces could apply a stateless filter designed to prevent devices in the home from originating source-address-spoofed traffic.

Using a filter configuration pseudo-syntax:

```
from !{interior_nets} to !{interior_nets} deny
... # permit or deny other kinds of traffic
```

5.3.2. Stateful filtering on external interfaces

Given knowledge of all interior network prefixes and the categorization of interfaces, all exterior interfaces could apply a stateful filter designed to discard traffic without matching state in the homenet router.

Using a filter configuration pseudo-syntax:

```
... # permit other kinds of good traffic first
from {interior_nets} to !{interior_nets} permit
from !{interior_nets} to {interior_nets} established permit
from any to any deny
```

5.3.3. Mixed-mode interface filtering

Given knowledge of all interior network prefixes and the categorization of interfaces, all mixed-mode interfaces could apply a stateful filter designed to discard exterior traffic without matching state in the homenet router and still statelessly permit internal traffic.

Using a filter configuration pseudo-syntax:

```
... # permit other kinds of good traffic first
from {interior_nets} to !{interior_nets} permit
from !{interior_nets} to {interior_nets} established permit
from {interior_nets} to {interior_nets} permit
from any to any deny
```

Because routing changes elsewhere in the home may cause traffic to shift among interior next hops which may not have state, traffic between interior routers may not be well-served by stateful filtering. One consequence for this policy on mixed-mode interfaces is that traffic from the exterior with spoofed source addresses from the "{interior_nets}" set of prefixes may be mistakenly allowed into the home.

Filter implementations which can incorporate knowledge of the previous and next hops and their classifications can design much more precise filters. Such implementations could deny traffic with "{interior_nets}" source addresses arriving from an exterior next hop, but permit them from an interior next hop on the same mixed-mode

interface.

6. Acknowledgements

Many thanks for the constructive input and criticism of Shwetha Bhandari, Lorenzo Colitti, Markus Stenberg, Mark Townsley, and Ole Troan.

Thanks also must go to pleasant, peaceful and productive trips on the Japan Rail (JR) Shinkansen ("bullet train").

7. IANA Considerations

This memo includes no request to IANA.

8. References

8.1. Normative References

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

8.2. Informative References

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007.

[RFC4890] Davies, E. and J. Mohacsi, "Recommendations for Filtering ICMPv6 Messages in Firewalls", RFC 4890, May 2007.

[RFC6092] Woodyatt, J., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, January 2011.

[TORPIG] Stone-Gross, B., "Your Botnet is My Botnet: Analysis of a Botnet Takeover", 2009, <<http://www.cs.ucsb.edu/~seclab/projects/torpig/torpig.pdf>>.

Author's Address

Erik Kline
Google Japan
Roppongi 6-10-1, 26th Floor
Minato, Tokyo 106-6126
JP

Phone: +81 03 6384 9000

Email: ek@google.com

