IPv6 Operations
Internet-Draft
Intended status: Informational
Expires: June 8, 2014

M. Gysi
Swisscom
G. Leclanche
Viagenie
E. Vyncke, Ed.
Cisco Systems
R. Anfinsen
Altibox

December 05, 2013

Balanced Security for IPv6 Residential CPE
draft-ietf-v6ops-balanced-ipv6-security-01

Abstract

This document describes how an IPv6 residential Customer Premise Equipment (CPE) can have a balanced security policy that allows for a mostly end-to-end connectivity while keeping the major threats outside of the home. It is documenting an existing IPv6 deployment by Swisscom and allows all packets inbound/outbound EXCEPT for some layer-4 ports where attacks and vulnerabilities (such as weak passwords) are well-known. The policy is a proposed set of rules that can be used as a default setting. The set of blocked inbound and outbound ports is expected to be updated as threats come and go.

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 June 8, 2014.

Table of Contents

1. Introduction

Internet access in residential IPv4 deployments generally consists of a single IPv4 address provided by the service provider for each home. The residential CPE then translates the single address into multiple private IPv4 addresses allowing more than one device in the home, but at the cost of losing end-to-end reachability. IPv6 allows all devices to have a globally unique IP address, restoring end-to-end reachability directly between any device. Such reachability is very powerful for ubiquitous global connectivity, and is often heralded as one of the significant advantages to IPv6 over IPv4. Despite this, concern about exposure to inbound packets from the IPv6 Internet (which would otherwise be dropped by the address translation function if they had been sent from the IPv4 Internet) remain.

This difference in residential default internet protection between IPv4 and IPv6 is a major concern to a sizable number of ISPs and the security policy described in this document addresses this concern without damaging IPv6 end-to-end connectivity.

The security model provided in this document is meant to be used as a pre-registered setting and potentially default one for IPv6 security in CPEs. The model departs from the "simple security" model described in [RFC6092] . It allows most traffic, including incoming unsolicited packets and connections, to traverse the CPE unless the CPE identifies the traffic as potentially harmful based on a set of rules. This policy has been deployed as a default setting in Switzerland by Swisscom for residential CPEs.

This document can be applicable to off-the-shelves CPE as well as to managed Service Provider CPE or for mobile Service Providers (where it can be centrally implemented).

2. Threats

For a typical residential network connected to the Internet over a broadband or mobile connection, the threats can be classified into:

- denial of service by packet flooding: overwhelming either the access bandwidth or the bandwidth of a slower link in the residential network (like a slow home automation network) or the CPU power of a slow IPv6 host (like networked thermostat or any other sensor type nodes);

- denial of service by Neighbor Discovery cache exhaustion [RFC6583]: the outside attacker floods the inside prefix(es) with packets with a random destination address forcing the CPE to exhaust its memory and its CPU in useless Neighbor Solicitations;

- denial of service by service requests: like sending print jobs from the Internet to an ink jet printer until the ink cartridge is empty or like filing some file server with junk data;

- unauthorized use of services: like accessing a webcam or a file server which are open to anonymous access within the residential network but should not be accessed from outside of the home

network or accessing to remote desktop or SSH with weak password protection;

- exploiting a vulnerability in the host in order to get access to data or to execute some arbitrary code in the attacked host;

- trojanized host (belonging to a Botnet) can communicate via a covert channel to its master and launch attacks to Internet targets.

3. Overview

The basic goal is to provide a pre-defined security policy which aims to block known harmful traffic and allow the rest, restoring as much of end-to-end communication as possible. This pre-defined policy should be centrally updated, as threats are changing over time. It could also be a member of a list of pre-defined security policies available to an end-customer, for example together with "simple security" from [RFC6092] and a "strict security" policy denying access to all unexpected input packets.

3.1. Rules for Balanced Security Policy

These are an example set of generic rules to be applied. Each would normally be configurable, either by the user directly or on behalf of the user by a subscription service. This document does not address the statefulness of the filtering rules as its main objective is to present an approach where some protocols (identified by layer-4 ports) are assumed weak or malevolent and therefore are blocked while all other protocols are assumed benevolent and are permitted.

If we name all nodes on the residential side of the CPE as 'inside' and all nodes on the Internet as 'outside', and any packet sent from outside to inside as being 'inbound' and 'outbound' in the other direction, then the behavior of the CPE is described by a small set or rules:

1. Rule RejectBogon: apply ingress filtering in both directions per [RFC3704] and [RFC2827] for example with unicast reverse path forwarding (uRPF) checks (anti-spoofing) for all inbound and outbound traffic (implicitly blocking link-local and ULA in the

same shot), as described in Section 2.1 Basic Sanitation and Section 3.1 Stateless Filters of [RFC6092];

2. Rule AllowManagement: if the CPE is managed by the SP, then allow the management protocols (SSH, SNMP, syslog, TR-069, IPfix, ...) from/to the SP Network Operation Center;

3. Rule ProtectWeakServices: drop all inbound and outbound packets whose layer-4 destination is part of a limited set (see Section 3.2), the intent is to protect against the most common unauthorized access and avoid propagation of worms; an advanced residential user should be able to modify this pre-defined list;

4. Rule Openess: allow all unsolicited inbound packets with rate limiting the initial packet of a new connection (such as TCP SYN, SCTP INIT or DCCP-request, not applicable to UDP) to provide very basic protection against SYN port and address scanning attacks. All transport protocols and all non-deprecated extension headers are accepted. This is a the major deviation from REC-11, REC-17 and REC-33 of [RFC6092].

5. All requirements of [RFC6092] except REC-11, REC-18 and REC-33 must be supported.

3.2. Rules Example for Layer-4 Protection: Swisscom Implementation

As of 2013, Swisscom has implemented the rule ProtectWeakService as described below. This is meant as an example and must not be followed blindly: each implementer has specific needs and requirements. Furthermore, the example below will not be updated as time passes, whereas threats will evolve.

Transport	Port	Description

	tcp	22	Secure Shell (SSH)
	tcp	23	Telnet
	tcp	80	HTTP
	tcp	3389	Microsoft Remote Desktop Protocol
	tcp	5900	VNC remote desktop protocol

Table 1: Drop Inbound

Transport	Port	Description
tcp-udp	88	Kerberos
tcp	111	SUN Remote Procedure Call
tcp	135	MS Remote Procedure Call
tcp	139	NetBIOS Session Service
tcp	445	Microsoft SMB Domain Server
tcp	513	Remote Login
tcp	514	Remote Shell
tcp	548	Apple Filing Protocol over TCP
tcp	631	Internet Printing Protocol
udp	1900	Simple Service Discovery Protocol
tcp	2869	Simple Service Discovery Protocol
udp	3702	Web Services Dynamic Discovery
udp	5353	Multicast DNS
udp	5355	Link-Lcl Mcast Name Resolution

Table 2: Drop Inbound and Outbound

Choosing services to protect is not an easy task, and as of 2013 there is no public service proposing a list of ports to use in such a policy. The Swisscom approach was to think in terms of services, by defining a list of services that are LAN-Only (ex: Multicast DNS) whose communication is denied by the policy both inbound and outbound, and a list of services that are known to be weak or vulnerable like management protocols that could be activated unbeknownst to the user.

The process used to set-up and later update the filters is out of scope of this document. The update of the specific rules could be

done together with a firmware upgrade or by a policy update (for example using Broadband Forum TR-069).

Among other sources, [DSHIELD] was used by Swisscom to set-up their filters. Another source of information could be the appendix A of [TR124]. The L4-filter as described does not block GRE tunnels ([RFC2473]) so this is a deviation from [RFC6092].

Note: the authors believe that with a dozen of rules only, a naive and unaware residential subscriber would be reasonably protected. Of course, technically-aware susbcribers should be able to open other applications (identified by their layer-4 ports or IP protocol numbers) through their CPE using some kind of user interface or even to select a completely different security policy such as the open or 'closed' policies defined by [RFC6092]. This is the case in the Swisscom deployment.

It is worth mentioning that PCP ([RFC6887]), UPnP ([IGD]) and similar protocols can also be used to dynamically override the default rules.

4. IANA Considerations

There are no extra IANA consideration for this document.

5. Security Considerations

The security policy protects from the following type of attacks:

- Unauthorized access because vulnerable ports are blocked

Depending on the extensivity of the filters, certain vulnerabilities could be protected or not. It does not preclude the need for end-devices to have proper host-protection as most of those devices

(smartphones, laptops, etc.) would anyway be exposed to completely unfiltered internet at some point of time. The policy addresses the major concerns related to the loss of stateful filtering imposed by IPV4 NAPT when enabling public globally reachable IPv6 in the home.

To the authors' knowledge, there has not been any incident related to this deployment in Swisscom network, and no customer complaints have been registered.

This set of rules cannot help with the following attacks:

- Flooding of the CPE access link;

- Malware which is fetched by inside hosts on a hostile web site (which is in 2013 the majority of infection sources).

6. Acknowledgements

The authors would like to thank several people who initiated the discussion on the ipv6-ops@lists.cluenet.de mailing list and others who provided us valuable feedback and comments, notably: Tore Anderson, Rajiv Asati, Fred Baker, Lorenzo Colitti, Paul Hoffman, Merike Kaeo, Simon Leinen, Eduard Metz, Martin Millnert, Benedikt Stockebrand. Thanks as well to the following SP that discussed with the authors about this technique: Altibox, Swisscom and Telenor.

7. Informative References

[DSHIELD] DShield, "Port report: DShield", <https://secure.dshield.org/portreport.html?sort=records>.

[IGD] UPnP Forum, "WANIPConnection:2 Service", December 20110, <http://upnp.org/specs/gw/UPnP-gw-WANIPConnection-v2-Service.pdf>.

[RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, December 1998.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

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

[RFC6092] Woodyatt, J., "Recommended Simple Security Capabilities in

Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, January 2011.

[RFC6583] Gashinsky, I., Jaeggli, J., and W. Kumari, "Operational Neighbor Discovery Problems", RFC 6583, March 2012.

[RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, April 2013.

[TR124] Broadband Forum, "Functional Requirements for Broadband Residential Gateway Devices", December 2006, <http://www.broadband-forum.org/technical/download/TR-124.pdf>.

Authors' Addresses

Martin Gysi
Swisscom
Binzring 17
Zuerich 8045
Switzerland

Phone: +41 58 223 57 24
Email: Martin.Gysi@swisscom.com

Guillaume Leclanche
Viagenie
246 Aberdeen
Quebec, QC G1R 2E1
Canada

Phone: +1 418 656 9254
Email: guillaume.leclanche@viagenie.ca

Eric Vyncke (editor)
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium

Phone: +32 2 778 4677
Email: evyncke@cisco.com

Ragnar Anfinsen
Altibox
Breiflaatveien 18
Stavanger 4069
Norway

Phone: +47 93488235
Email: Ragnar.Anfinsen@altibox.no