Advanced Security for IPv6 CPE
draft-vyncke-advanced-ipv6-security-03.txt

Abstract

This document describes how an IPv6 residential Customer Premise
Equipment (CPE) can leverage modern security techniques to have strong
security, while retaining as much of the end-to-end reachability of
IPv6 as possible.
It is a re-submission in the framework of the HOMENET working group.
The reputation part of this document should leverage the work done in
the REPUTE working group of the Application are.
The IETF has been notified of intellectual property rights claimed in
regard to some or all of the specification contained in this document.
For more information consult the online list of claimed rights.

Status of this Memo

Copyright Notice

1. Introduction

Internet access in residential IPv4 deployments generally consist of a single IPv4 address provided by the service provider for each home. Residential CPE then translates the single address into multiple private 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 unique, global, 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 document describes firewall functionality for an IPv6 CPE which departs from the "simple security" model described in [RFC6092] . The intention is to provide an example of a security model which 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 signatures (and other correlation data and heuristics) that are kept up to date on a regular basis. The computational resources necessary to support some, not all, functionalities of this model are likely more intensive than those described in [RFC6092], but are easily within the realm of what is commonly available in 2011 on medium to

high-end network based firewall systems for small and medium businesses, or host-based commercial firewalls that run on laptop and desktop PCs. This set of techniques is also known as Universal Threat Mitigation (UTM).

2. Threats

For a typical residential network connected to the Internet over a broadband 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 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 freely and anonymously from outside of the home network

*exploiting a vulnerability in the host in order to get access to data or to execute some arbitrary code in the attacked host. Exploitation can be further divided in two classes:

   1. day-0 attack when this attack has never been seen before (hence nothing can really detect it) and

   2. day+n attack where this attack is known and can be detected by the use of an attack signature

*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 an adaptive security policy which aims to block known harmful traffic and allow the rest, restoring as much of end-to-end communication as possible. In addition, new protocols may evolve and be deployed over time; only if they become a threat vector does the CPE firewall receive a signature update (including dynamic correlation data) to classify and block them. This is in direct contrast to [RFC6092], which requires built-in knowledge of a number of protocols, or requires Internet communication to be limited to a handful of protocols that the CPE understands how to process.

*Intrusion Prevention System (IPS) is a signature-based technology
 which inspects a pre-defined set of protocols at all layers (from
 layer-3 to layer-7) and uses a vast set of heuristics to detect
 attacks within one or several flow. Upon detection, the flow is
 terminated and an event is logged for further optional auditing.
 As exploits are added every day, the signature database must be
 updated daily and is usually quite large (more than 100 MB). This
 requires both large local storage (large flash or even a hard
 disk) and a subscription to an update service.

*Reputation database is a centralized database which gives a
 reputation score to any IPv6 address (or prefix). The score
 varies from untrusted to trusted. Untrusted IPv6 addresses are
 typically addresses of a well-known attacker or from a Botnet
 member or from an ISP with a poor track of security... Protocols
 exist to dynamically request a reputation (based on DNS or HTTP).
 This usually requires a subscription. Note: in IPv6 the
 reputation database concept is still in its infancy, for example,
 little experience exists on the scope of the reputation: a host /
 128, a LAN prefix /64 or a delegated prefix size of /56 or /48...

*Local correlation uses another set of heuristics (like TCP
 distribution of Initial Sequence Number or used TCP ports or
 protocol handshake banners) to assert the variety of local hosts
 (namely operating system (OS) version and set of application) and
 raise or decrease the importance of a specific attack signature.
 For example, if the OS of host A is OS-A, then there is no point
 to inspect traffic to or from host A for attacks which are only
 relevant to OS-B.

*Global correlation leverage all IPS distributed on the Internet
 to build the reputation database as well as changing the
 relevance of an IPS signature (for example, a propagating worm
 will trigger a lot of identical signatures on several IPS, this
 should raise the relevance of a specific signature up to the
 point of blocking all inbound/outbound connections on a specific
 layer-4 port).

The above techniques are common in the large network where budget is
enough to buy firewalls, IPS and subscribe to signature or reputation
source. The authors of this document believes that competition and
Moore's law will make the set of those techniques (commonly referred to
as 'Universal Threat Mitigation') affordable for consumer space.

3.1. Rules for Security Policy

These are an example set of rules to be applied. Each would normally be
configurable, either by the user directly or on behalf of the user by a

subscription service. The default preferred state hasn't been listed,
though it is expected that all rules would be on by default.
If we named all hosts on the residential side of the CPE as 'inside'
and all hosts on the Internet as 'outside', then the behavior of the
CPE is described by a small set or rules:

1. Rule RejectBogon: apply unicast reverse path forwarding (RPF)
   checks (anti-spoofing) for all inbound and outbound traffic
   (implicitly blocking link-local and ULA in the same shot)

2. Rule BlockBadReputation: block all inbound and outbound packets
   whose outside IPv6 address has a bad reputation score

3. Rule AllowReturn: inspect all outbound traffic and allow the
   return traffic matching the states (5-tuple + TCP sequence
   number or any layer-4 state), apply IPS on the outbound (to
   block Botnet) and inbound (to block malicious/cracked servers
   which could inject malware) with IPS. If the protocol is not
   supported/recognized by the IPS, accept it anyway.

4. Rule AllowToPublicDnsHost: allow all inbound traffic to any
   inside address which is listed in the public DNS with a AAAA
   record (this requires that the CPE/RG can do a zone transfer,
   i.e., that the CPE/RG appears like a secondary name server),
   all inbound traffic is also inspected with IPS. If the protocol
   is not supported/recognized by the IPS, accept it anyway.

5. Rule ProtectLocalOnly: block all inbound traffic to any inside
   address as long as the inside address has never sent a packet
   to the outside. The intent is to protect local-only devices
   like thermostat or printers. Most (if not all) hosts expecting
   inbound connections have to send a couple of outbound packets
   to the outside (registration, DNS request, ...). This is the
   usual IPv4 firewall behavior augmented with IPS and reputation

6. Rule CrypoIntercept: at the exception of IPsec, all inbound
   connections that are encrypted (notably TLS [RFC5246]) must be
   intercepted (this is terminated by the CPE that will present
   its own self-signed certificate to the remote party which
   should have installed the CPE self-signed certificate in a
   secure way in its trust anchors store) in order to allow for
   further inspection. The decrypted flow is then passed again
   through those rules and encrypted again before being forwarded
   to the local host. This is actually a Man-in-the-Middle attack
   done for a good reason: protect the naive residential user. Of
   course, documentation and GUI MUST be provided to educate the
   user and help him/her to understand how to do it in a secure
   way. Note: this technique is also used nowadays by large

enterprise web proxies with the self-signed certificate being
securely distributed to all clients.

7. Rule ParanoidOpeness: allow all unsolicited inbound connections
rate limited to protect against port and address scanning
attacks or overloading devices or slow links within the home.
The connection MUST be inspected by the IPS engine. If the
connection is anonymous or using a default password (like
connecting to a webcam as a guest), then the flow SHOULD be
dropped. If the IPS detects an attack, then the flow MUST be
closed. If the protocol is not recognized as supported by the
IPS, the flow MAY be allowed.

3.2. Security Analysis

This proposal of 'paranoid openness' stops the following attacks:

*unauthorized use of services/denial of service: because all
 anonymous access to inside servers are blocked.

*Denial of services on low bandwidth or low CPU inside hosts IFF
 those hosts never access the Internet

*Exploiting of a day+1 attack, those attacks are blocked with the
 IPS signature and address reputation database

The CryptoIntercept part can also be leveraged as a small Certification
Authority (CA) that could generate RSA key pairs and X.509 certificates
at the CPE/RG owner's request. Those key pairs and certificates can
then be given to trusted devices or users (like the owner's laptop so
that he/she could easily and safely connect from the outside).
This proposal cannot help with the following attacks:

*flooding the access link to the Internet, this is exactly the
 same as with the old layers-3/4 firewall approach as only the ISP
 can effectively stop the flooding of the CE-PE link;

*weak password on inside services, of course the IPS component
 will detect multiple failed attempts (dictionary attack) and
 report the offender to the Global Correlation system;

*exploiting of day-0 attack: until now, these day-0 attacks are
 caused either by rapidly propagating worms (then the global
 correlation of unusual traffic pattern will raise an alert and
 block the traffic after a couple of hundred's of successful
 attacks) or by targeted attacks against high-profile targets
 (like Government or banks or ;..) which should be protected by
 conventional less open security policies;

     *exploiting a vulnerability in a rare or new protocol (not yet
      supported by the IPS), this case will probably never occur on a
      wide scale in a residential use of Internet.

4. IANA Considerations

There are no extra IANA consideration for this document.

5. Security Considerations

All security considerations have been done in the Security Analysis
[Section 3.2](#).
It is also advisable that the inbound rate limiter system could be
added to the [RFC6092] as it is light and does not depend on a
centralized policy server.

6. Acknowledgements

Many thanks to Ole Troan, Stuart Cheshire, Dave Oran and Eliot Lear for
the review of the -00 version and to Ron Bonica, Sam Hartmans, Lee
Howard, Greg Lebovitz, Jordi Palet, Tina Tsou and others for their
comments during and after the first presentation at the Hiroshima IETF
meeting in November 2009.
A previous IETF work has similar ideas [I-D.palet-v6ops-ipv6security].

7. References

7.1. Normative References

[RFC5246]	Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

7.2. Informative References

[RFC2993]	Hain, T., "Architectural Implications of NAT", RFC 2993, November 2000.
[RFC6092]	Woodyatt, J., "Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service", RFC 6092, January 2011.
[I-D.palet-v6ops-ipv6security]	Palet, J, Vives, A, Martinez, G and A Gomez, "IPv6 distributed security requirements", Internet-Draft draft-palet-v6ops-ipv6security-02, February 2005.

Authors' Addresses

   Eric Vyncke Vyncke Cisco Systems De Kleetlaan 6a Diegem, 1831
   Belgium Phone: +32 2 778 4677 EMail: evyncke@cisco.com

Andrew Yourtchenko Yourtchenko Cisco Systems De Kleetlaan 6a
Diegem, 1831 Belgium Phone: +32 2 704 5494 EMail: ayourtch@cisco.com

Mark Townsley Townsley Cisco Systems 11, Rue Camille Desmoulins Issy
Les Moulineaux, 92782 France Phone: +33 15 804 3483 EMail:
townsley@cisco.com