

6man Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 17, 2012

T. Macaulay
Bell Canada
February 14, 2012

IPv6 packet staining
draft-macaulay-6man-packet-stain-00

Abstract

This document specifies the application of security staining on an IPv6 datagrams and the minimum requirements for IPv6 nodes staining flows, IPv6 nodes forwarding stained packets and interpreting stains on flows.

The usage of the packet staining destination option enables proactive delivery of security intelligence to IPv6 nodes such as firewalls and intrusion prevention systems, and end-points such servers, workstations, mobile and smart devices and an infinite array of as-yet-to-be-invented sensors and controllers.

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 August 17, 2012.

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
2. Conventions used in this document	3
3. Background	3
3.1. Packet Staining Benefits	4
3.2. Implementation and support models	5
3.3. Use cases	5
4. Requirements for staining IPv6 packets	6
5. Packet Stain Destination Option (PSDO)	7
6. Acknowledgements	8
7. Security Considerations	8
8. IANA Considerations	9
9. Normative References	9
Author's Address	10

1. Introduction

From the viewpoint of the network layer, a flow is a sequence of packets sent from a particular source to a particular unicast, anycast, or multicast destination. From an upper layer viewpoint, a flow could consist of all packets in one direction of a specific transport connection or media stream. However, a flow is not necessarily 1:1 mapped to a transport connection.

Traditionally, flow classifiers have been based on the 5-tuple of the source and destination addresses, ports, and the transport protocol type. However, as the growth of internetworked devices continues under IPv6, security issues associated with the reputation of the source of flows are becoming a critical criterion associated with the trust of the data payloads and the security of the destination endpoints and the networks on which they reside.

The usage of security reputational intelligence associated with the source address field and possibly the port and protocol [REF1] enables packet-by-packet IPv6 security classification, where the IPv6 header extensions in the form of Destination Options may be used to stain each packet with security reputation information such that the network routing is unaffected, but intermediate security nodes and endpoint devices can apply policy decisions about incoming information flows without the requirement to assemble and treat payloads at higher levels of the stack.

IPv6 packet staining support consists of labeling datagrams with security reputation information through the addition of an IPv6 destination option in the packet header by packet manipulation devices (PMDs) in the carrier or enterprise network. This destination option may be read by in-line security nodes upstream from the packet destination, as well as by the destination nodes themselves.

2. Conventions used in this document

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 [RFC2119].

3. Background

Internet based threats in the form of both malicious software and the agents that control this software (organized crime, spys, hacktivits) have surpassed the abilities of signature-based security

systems; whether they be on the enterprise perimeter, within the corporate network, on the endpoint point or in-the-cloud (internet-based service). Additionally, the sensitivity of IP network continues to grow as new generation of smart devices is appearing on the networks in the form of broadband mobile devices, legacy industrial control devices, and very low-power sensors. This diverse collections of IP-based assets is coming to be known as the Internet of Things (IOT).

In response to the accelerating threats, the security vendor community have integrated their products with proprietary forms of security reputation intelligence. This intelligence is about IP addresses and domains which have been observed engaged in attack-behaviours such as inappropriate messaging and traffic volumes, domain management, Botnet command-and-control channel exchanges and other indicators of either compromise or malicious intent. [REF 1] IP address may also end up on a security reputation list if they are identified as compromised through vendor-specific signature-based processes. Security reputation intelligence from vendors is typically made available to perimeter and end-point products through proprietary, internet-based queries to vendor information bases.

This system of using proactive, security reputation intelligence has many benefits, but also several weakness and scaling challenges. Specifically, existing intelligence systems are:

1. subject to direct attack from the internet on distribution points, for instance
2. are proprietary to vendor devices
3. require fat-clients consuming both bandwidth and CPU, and
4. introduces flow latency while queries are sent, received and processed
5. introduces intelligence latency as reputation lists will be inevitably cached and only periodically refreshed given the number and range of vendor-specific processing elements

3.1. Packet Staining Benefits

In contrast to the challenges of current security reputation intelligence systems, packet staining has the following strengths

1. packet staining can occur transparently in the network, presenting no attack surface
2. packet staining uses standardized, public domain IPv6 capabilities
3. security rules can be easily applied in hardware or firmware
4. reading packet stains introduces little to no latency
5. near-real-time threat intelligence distribution systems can be implemented can be implemented out of band in PMDs using a standardized packet staining method allowing multiple

intelligence sources (vendor sources) to be aggregated and applied in an agnostic (cross-vendor) manner.

3.2. Implementation and support models

Packet staining may be accomplished by different entities including carriers, enterprises and third-party value-added service providers.

Carriers or service providers may elect to implement staining centres at strategic locations in the network to provide value-added services on a subscription basis. Under this model, subscribers to a security staining service would see their traffic directed through a staining centre where Destination Options are added to the IPv6 headers and IPv4 traffic is encapsulated within IPv6 tunnels, with stained headers.

Carriers or service providers may elect to stain all IPv6 traffic entering their network, and allow subscribers to process the stains at their own discretion.

If such upstream, network-based staining services are inappropriate or unavailable, Enterprise data centre managers / cloud computing service providers may elect to deploy IPv6 staining at the perimeter into the internal network, tunnelling all IPv4 traffic, and allow data centre/cloud service users to process stains at their discretion.

Enterprise may wish to deploy IPv6 on internal networks, and stain all internal traffic whereby security nodes and end-points may apply corporate security policy related to reputation.

3.3. Use cases

The following are example use-cases for a security technique based upon a packet staining system.

Organization Perimeter Use-case Traffic to a subscriber is routed through a PMD in the carrier network configured to stain (apply Destination Options extensions) all packets to the subscriber (TM)s IP-range, which have entries in the threat intelligence information base. The PMD accesses the information base from a locally cached file or other method not defined in this draft. Packets from sources not in the information base pass through the PDM unchanged. Packets from sources in the information base have a Destinations Option added to the datagram header. The Destination Options contains reputation from the information base. The format of the destination option is discussed later in this draft. IPv6 perimeter devices such as firewalls, web proxies or security routers on the perimeter of the

subscriber network look for Destination Options on incoming packets with reputation stains. If a stain is found, the perimeter device applies the organization policy associated with the reputation indicated by the stain. For instance, drop the packet, quarantine the packet, issue alarms, or pass the packets and associated flow to specially hardened extra-net authentication systems, or do nothing.

IPv4 support Use-case" IPv4 header fields and options are not suitable for packet staining; however, there is a clear security benefit to supporting IPv4 flows. IPv4 traffic to a subscriber is routed through a PMD in the carrier network configured to encapsulate the IPv4 traffic in an IPv6 tunnel. The PMD applies a stain (Destination Options extension) to the IPv6 tunnel as per the Perimeter Use-case above. Subscriber perimeter devices such as firewalls, web proxies or security routers are configured to support both native IPv6 flows and IPv6 tunnels contain legacy IPv4 flows. Perimeter devices look for Destination Options on incoming IPv6 packets with reputation stains. If a stain is found, the perimeter device applies the organization policy associated with the reputation indicated by the stain to the IPv4 packet within the IPv6 tunnel. In this manner IPv4 support may be transparent to end-users and applications.

IPv6 end-point use-case" IPv6 end-points may make use of reputation stains by processing Destination Options before engaging in any application level processing. In the case of certain classes of smart device, remote and mobile sensors, reputation stains may be a critical form of security when other mitigations such as signature bases and firewalls are too power and processor intensive to support.

URL-specific stains" it is a common occurrence to see large public content portals with millions of users sharing dozens of addresses. Frequently, malicious content will be loaded to such sites. This content represents a very small fraction of the otherwise legitimate content on the site, which may be under the direct control of entirely separate entities . Degrading the reputation of IP addresses used by these large portals based on a very small amount of content is problematic. For such sites, reputation stains should have the ability to include the URL of malicious content, such that the reputation of the only specific portions of these large portals is degraded according to threat evidence, rather than the entire IP address, CIDR block, ASN or domain name.

4. Requirements for staining IPv6 packets

1. The default behaviour of a security node MUST be to leave a packet unchanged (apply no stain).
2. Reputation stains may be inserted or overwritten by security nodes in the path.
3. Reputation stains may not be applied by the sender/source of the packet.
4. The reputation staining mechanism needs to be visible to all stain-aware nodes on the path.
5. The mechanism needs to be able to traverse nodes that do not understand the reputation stains. This is required to ensure that packet-staining can be incrementally deployed over the Internet.
6. The presence of the reputation staining mechanism should not significantly alter the processing of the packet by nodes, unless policy is explicitly configured. This is required to ensure that stained packets do not face any undue delays or drops due to a badly chosen mechanism.
7. A PMD should be able to distinguish a trusted stain from an untrusted stain, through mechanism such as digital signatures or intrinsic trust among network elements.
8. A staining node MAY apply more specific and selective staining services according to subscriptions. Staining nodes SHOULD support different reputation taxonomies to support different subscribers and/or interoperability with other staining entities, and have the ability to stain flows to different subscriber sources according to different semantics.

5. Packet Stain Destination Option (PSDO)

The Packet Stain Destination Option (PSDO) is a destination option that can be included in IPv6 datagrams that are inserted by PMDs in order to inform packet staining aware nodes on the path, or endpoints, that the PSDO has an alignment requirement of (none).

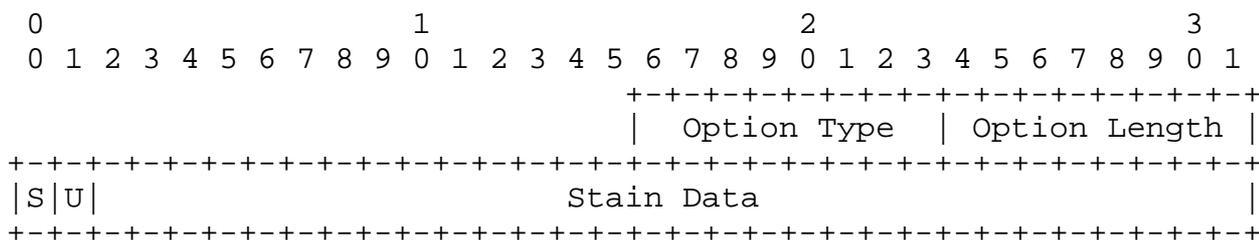


Figure 1: Packet Stain Destination Option Layout

Option Type

8-bit identifier of the type of option. The option identifier for the reputation stain option will be allocated by the IANA.

Option Length

8-bit unsigned integer. The length of the option (excluding the Option Type and Option Length fields).

S Bit

When this bit is set, the reputation stain option has been signed.

U Bit

When this bit is set, the reputation stain option contains a malicious URL.

Stain Data

Contains the staining data.

6. Acknowledgements

The author wishes to acknowledge the guidance and support of Suresh Krishnan from Ericsson's Montreal lab. The author also wishes to credit Chris Mac-Stoker from NIKSUN for his substantial contributions to the early stages of the packet staining concept.

7. Security Considerations

Some implementation may elect to not apply digital signature to reputation stains in the Destination Option, in which case the stain is not protected in any way, even if IPsec authentication [RFC4302] is in use. Therefore an unsigned reputation stain can be forged by an on-path attacker. Implementers are advised that any en-route change to an unsigned security reputation stain value is undetectable. Therefore packet staining use the Destination Options extension without digital signatures requires intrinsic trust among the network elements and the PMD, and the destination node or intervening security nodes such as firewalls or IDS services. For this reason, receiving nodes MAY need to take account of the network from which the stained packet was received. For instance, a multi-homed organization may have some service providers with staining

services and others that do not. A receiving node SHOULD be able to distinguish which source from which stains are expected. A receiving node SHOULD by default ignore any reputation stains from sources (networks or devices) that have not been specifically configured as trusted.

The reputation intelligence of IP source addresses, ASNs, CIDR blocks and domains is fundamental to the application of reputation stains within packet headers. Such reputation information can be seeded from a variety of open and closed sources. Poorly managed or compromised intelligence information bases can result in denial of service against legitimate IP addresses, and allow malicious entities to appear trustworthy. Intelligence information bases themselves may be compromised in a variety of ways; for instance the raw information feeds may be corrupted with erroneous information, alternately the intelligence reputation algorithms could be flawed in design or corrupted such that they generate false reputation scores. Therefore seed intelligence SHOULD be sourced and monitored with demonstratable diligence. Similarly, reputation algorithms should be protected from unauthorized change with multi-layered access controls.

The value of reputation stains will be directly proportional to the trustworthiness, reliability and reputation of the intelligence source itself. Operators of security nodes SHOULD have defined and auditable methods upon which they select and manage the source of reputation intelligence and the packet staining infrastructure itself.

8. IANA Considerations

This document defines a new IPv6 destination option for carrying security reputation packet stains. IANA is requested to assign a new destination option type (TBA1) in the Destination Options registry maintained at <http://www.iana.org/assignments/ipv6-parameters> 1) Signed Security Reputation Option, 2) Unsigned Security Reputation Option 3) Signed Security Reputation Option with malicious URL 4) Unsigned Security Reputation Option with malicious URL The act bits for this option need to be 10 and the chg bit needs to be 0.

9. Normative References

[REF1] Macaulay, T., "Upstream Intelligence: anatomy, architecture, case studies and use-cases.", Information Assurance Newsletter, DOD , Aug to February 2010 to 2011.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate

Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

Author's Address

Tyson Macaulay
Bell Canada
160 Elgin Floor 5
Ottawa, Ontario
Canada

Email: tyson.macaulay@bell.ca