

IP Multicast  
Internet-Draft  
Intended status: Informational  
Expires: June 26, 2016

E. Vyncke  
Cisco  
E. Rey  
ERNW  
A. Atlasis  
NCI Agency  
December 24, 2015

**MLD Security**  
**draft-vyncke-pim-mld-security-01**

**Abstract**

The latest version of Multicast Listener Discovery protocol is defined in [RFC 3810](#), dated back in 2004, while the first version of MLD, which is still in use and has not been deprecated, is defined in [RFC 2710](#) and is dated back in 1999. New security research has exhibited new vulnerabilities in MLD, both remote and local attack vectors. This document describes those vulnerabilities and proposes specific mitigation techniques.

**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 26, 2016.

**Copyright Notice**

Copyright (c) 2015 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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Local Vulnerabilities . . . . .	<a href="#">3</a>
<a href="#">2.1.</a>	Downgrading to MLDv1 . . . . .	<a href="#">3</a>
<a href="#">2.2.</a>	Queries sent to unicast address . . . . .	<a href="#">4</a>
<a href="#">2.3.</a>	Win the election . . . . .	<a href="#">4</a>
<a href="#">2.4.</a>	Host enumeration and OS fingerprinting . . . . .	<a href="#">4</a>
<a href="#">2.5.</a>	Flooding of MLD messages . . . . .	<a href="#">4</a>
<a href="#">2.6.</a>	Amplification . . . . .	<a href="#">4</a>
<a href="#">3.</a>	Remote Vulnerabilities . . . . .	<a href="#">5</a>
<a href="#">4.</a>	Mitigations . . . . .	<a href="#">5</a>
<a href="#">5.</a>	IANA Considerations . . . . .	<a href="#">5</a>
<a href="#">6.</a>	Security Considerations . . . . .	<a href="#">6</a>
<a href="#">7.</a>	Acknowledgements . . . . .	<a href="#">6</a>
<a href="#">8.</a>	References . . . . .	<a href="#">6</a>
<a href="#">8.1.</a>	Normative References . . . . .	<a href="#">6</a>
<a href="#">8.2.</a>	Informative References . . . . .	<a href="#">6</a>
	Authors' Addresses . . . . .	<a href="#">7</a>

## [1.](#) Introduction

The Multicast Listener Discovery protocol version 2 (MLDv2) [RFC3810](#) [[RFC3810](#)] has a security section but it was not exhaustive and the focus was only on local forged MLD packets. The same is also true for the first version of MLD (now called MLDv1), which is still in use, defined in [RFC 2710](#). This document goes beyond those attacks.

For the reader who is not familiar with MLDv2, here are the main points:

Multicast routers send MLD queries which are either generic (query about all multicast group) sent to ff02::1 (link-scope all nodes) or specific (query about a specific group) sent to this multicast group. Query messages can also be sent to a unicast address.

Multicast members reply to MLDv2 queries with reports sent to ff02::16 (link-scope all MLDv2 routers). In version 1 of MLD [RFC2710](#) [[RFC2710](#)], the reports are sent to the multicast group being reported. Reports can be transmitted twice or more in order to ensure that the MLD router gets at least one report.



When a node ceases to listen to a multicast address on an interface, it sends an MLDv1 Done message or a specially crafted MLDv2 Report message.

All MLD packets are ICMPv6 [RFC4443](#) [[RFC4443](#)] messages sent with a hop-limit of 1, from a link-local address and there is no authentication.

MLD messages received with a hop-limit greater than 1 should be discarded.

Neighbor Discovery Protocol [RFC4861](#) [[RFC4861](#)] requires nodes to become member of the respective solicited-node multicast groups for all their link-scope and global-scope addresses.

Switches are assumed to implement MLD snooping [RFC4541](#) [[RFC4541](#)] to learn where to forward multicast packets. It must be noted though that implementations of MLD snooping do not act on link-local multicast groups such as solicited-node multicast group: they simply forward all packets destined to a link-local multicast group to all port in the same layer-2 network.

MLDv2 was designed to be interoperable with MLDv1.

The main difference between MLDv1 and MLDv2 from a functionality perspective is that MLDv1 does not support "source filtering" (in MLDv2 nodes can report interest in traffic only from a set of source addresses or from all except a set source addresses).

Every IPv6 node must support MLD.

This document is heavily based on previous research: [[Troopers2015](#)].

## **2. Local Vulnerabilities**

### **[2.1.](#) Downgrading to MLDv1**

A single MLDv1 report message is enough to downgrade all MLD nodes (hosts and routers) to the version 1 protocol. This could be used to force a MLD host to reply with MLDv1 reports sent to the multicast group rather than to ff02::16. This downgrade to MLDv1 could also be used to transmit the MLDv1 report with a 'done' operation to remove a listener (stopping the multicast traffic on the subnet). Another consequence of downgrading to MLDv1 can be the fact that an attacker can also used "Host Suppression" feature as part of a DoS attack, make the launch of such an attack easier.



## **2.2. Queries sent to unicast address**

[Section 5.1.15 of RFC3810](#) [[RFC3810](#)], specifies that for debugging purposes, nodes must accept and process queries sent to any of their addresses (including unicast). Lab testing, described in [[Troopers2015](#)], clearly shows that all implementations except FreeBSD accept and process MLD queries sent to a unicast global address. This can be exploited to completely bypass the legitimate MLD router and interact directly (for whatever purpose) with the targets (including legitimate routers and clients).

## **2.3. Win the election**

When there are multiple MLD routers in a layer-2 domain, the one with the lowest IPv6 address wins the election and becomes the designated MLD router. A hostile node can then send from a lower link-local address an MLD message and become the MLD router. This fact in combination with the direct interaction with the targets could be leveraged to mount a denial of service attack.

## **2.4. Host enumeration and OS fingerprinting**

Some hosts try to prevent host enumeration by not responding to ICMPv6 echo request messages sent to any multicast group. But, the same hosts must reply to any MLD queries including the generic one sent to ff02::1, this allows for MLD host enumeration. As hosts join different groups based on their operating system (specific groups for Microsoft Windows for example), the MLD report can also help for Operating System (OS) fingerprinting.

## **2.5. Flooding of MLD messages**

If an implementation does not rate limit in hardware the rate of processed MLD messages, then they are vulnerable to a CPU exhaustion denial of services. If a node does not limit the number of states associated to MLD, then this node is vulnerable to a memory exhaustion denial of services.

## **2.6. Amplification**

Nodes usually join multiple groups (for example, Microsoft Windows 8.1 joins 4 groups). Therefore a forged generic MLDv1 query will force those nodes to transmit MLDv1 reports for each of their groups (in our example 4); furthermore, many implementations send MLD reports twice (in our example 8 in total). MLDv2 is a little better because reports are sent to ff02::16 and not to the multicast group.



### **3. Remote Vulnerabilities**

MLD messages with hop-limit different than 1 should be discarded but nothing prevents a hostile party located  $n$  hops away from the victim to send any MLD messages with a hop-limit set to  $n+1$ . Therefore, a remote hostile party can mount attacks against MLD (especially because implementations process MLD queries sent to a global unicast address).

### **4. Mitigations**

This section proposes some mitigation techniques that could be used to prevent the above attacks. This section is not a specification of any kind, the words 'should' is plain English and is not related to [RFC2119](#) [[RFC2119](#)].

Mitigation by specific implementations:

Similar to RA-guard [RFC6105](#) [[RFC6105](#)], there should be a MLD-guard function in layer-2 switches; MLD queries (either version 1 or version 2) received on ports attached to non multicast routers should be discarded. Switches could also block all MLDv1 packets in order to prevent the downgrading of MLD version. Of course, this requires all nodes to support MLDv2.

All nodes should be able to disable MLDv1.

Control plane policing should also be implemented in order to avoid denial of services attacks.

Mitigation by a protocol update of [RFC2710](#) [[RFC2710](#)] and [RFC3810](#) [[RFC3810](#)]:

MLD queries should not be accepted and processed when sent to a unicast address (either link-local or global scope). This requires update of [RFC 3810](#) and [RFC 2710](#).

To mitigate the remote attacks, the hop-limit should have been set to 255 and MLD nodes should discard packets with a hop-limit different than 255.

### **5. IANA Considerations**

This document contains no IANA considerations.



## **6. Security Considerations**

This document describes multiple vulnerabilities that have been described above and tries to mitigate them or even eliminate some of them by making specific suggestions for update of the protocol as well as by suggesting the implementation of related security mechanisms to layer-2 devices.

## **7. Acknowledgements**

The authors would like to thank Stig Venaas for some discussions on this topic.

## **8. References**

### **8.1. Normative References**

- [RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", [RFC 2710](#), DOI 10.17487/RFC2710, October 1999, <<http://www.rfc-editor.org/info/rfc2710>>.
- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", [RFC 3810](#), DOI 10.17487/RFC3810, June 2004, <<http://www.rfc-editor.org/info/rfc3810>>.

### **8.2. Informative References**

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 4443](#), DOI 10.17487/RFC4443, March 2006, <<http://www.rfc-editor.org/info/rfc4443>>.
- [RFC4541] Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", [RFC 4541](#), DOI 10.17487/RFC4541, May 2006, <<http://www.rfc-editor.org/info/rfc4541>>.



- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), DOI 10.17487/RFC4861, September 2007, <<http://www.rfc-editor.org/info/rfc4861>>.
- [RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", [RFC 6105](#), DOI 10.17487/RFC6105, February 2011, <<http://www.rfc-editor.org/info/rfc6105>>.
- [Troopers2015] Rey, E., Atlasis, A., and J. Salazar, "MLD Considered Harmful", 2015, <[https://www.troopers.de/media/filer\\_public/7c/35/7c35967a-d0d4-46fb-8a3b-4c16df37ce59/troopers15\\_ipv6secsummit\\_atlasis\\_rey\\_salazar\\_mld\\_considered\\_harmful\\_final.pdf](https://www.troopers.de/media/filer_public/7c/35/7c35967a-d0d4-46fb-8a3b-4c16df37ce59/troopers15_ipv6secsummit_atlasis_rey_salazar_mld_considered_harmful_final.pdf)>.

#### Authors' Addresses

Eric Vyncke  
Cisco  
De Kleetlaan 6a  
Diegem 1831  
Belgium

Phone: +32 2 778 4677  
Email: [evyncke@cisco.com](mailto:evyncke@cisco.com)

Enno Rey  
ERNW  
Carl-Bosch-Str. 4  
Heidelberg 69115  
Germany

Phone: +49 6221 480390  
Email: [erey@ernw.de](mailto:erey@ernw.de)

Antonios Atlasis  
NCI Agency  
Oude Waalsdorperweg 61  
The Hague 2597 AK  
The Netherlands

Phone: +31 703743564  
Email: [antonios.atlasis@ncia.nato.int](mailto:antonios.atlasis@ncia.nato.int)

