IPv6 Operations
Internet-Draft
Intended status: Informational
Expires: December 28, 2015

E. Vyncke
A. Yourtchenko
Cisco
D. Valenkamp
ETH Zurich

June 26, 2015

IPv6-Only for Wired Thin-Clients
draft-vyncke-v6ops-ipv6-only-thin-clients-00

Abstract

While IPv6-only (no IPv4 at all) is becoming an objective, there are remaining issues on this road for the wired thin clients. This document enumerates of couple of them; each with a short description, followed by a description of the issue in IPv6-only networks then some solutions.

It is expected that this document will grow by collecting other roadblocks and suggestions to remove them.

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 December 28, 2015.

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

1. Wake-on-Lan

1.1. Description

Wake-on-Lan also known as WOL is specified in [WOL]. It allows for a remote operator to wake a sleeping host in order to trigger software update while the host is sleeping (for example during the night of the week-end). It consists of sending a specially formatted frame for a specific host. This is called the magic packet: with the Ethernet payload having somewhere 6 bytes containing 0xFF followed by 16 times the network interface datalink-layer address.

1.2. Issue

As the host is sleeping, it does not transmit any packets and will not reply to neither ARP request nor Neighbor Solicitations. This means that the adjacent routers have lost the binding between datalink and network address and also that all layer-2 switches have lost the binding between the datalink-layer address and the port/ interface. So, it is not possible to send a unicast IPv4 or IPv6 packet containing this magic packet as it will be dropped by the router (no adjacency information). In IPv4, a local configuration can allow the 'directed broadcast' (see RFC2644 [RFC2644])such that

the magic packet can be sent to an IPv4 directed broadcast which will be sent to a datalink-layer broadcast, i.e. forwarded on all ports of all routers and switches in the same layer-2 domain. Therefore, the magic packet will reach all hosts including the sleeping one.

In IPv6, there is no directed broadcast for good reason. Only a link-local multicast group such as ff02::1 for all link-local hosts. So, the magic packet for a single host could be sent to this multicast group, reaching all link-local hosts (as switches and routers will forward this packet to all ports/interfaces) and waking up the sleeping node. But, there is no solution for a remote operator to send this magic packet...

1.3. Mitigation

A trivial solution would be to hard code in the router configuration a specific global or ULA address to the broadcast data-link layer address. For example, to reach all nodes in 2001:db8::/64, let’s configure a static Neighbor Cache entry for 2001:db8::cafe:c0:ffee as ff-ff-ff-ff-ff-ff. Then a remote operator can send the magic packet to this destination, it will be routed across the layer-3 network, will be addressed to the data-link layer broadcast address which will be flooded by all layer-2 switches on all their ports, finally reaching the sleeping host.

This approach has two drawbacks:

1. provisioning of all those mappings in all routers

2. opening a door to a denial of service attack: a remote hostile party could keep sending packets this is specific unicast address forcing all hosts to stay awake, hence wasting electrical energy. As this address is a unicast address which does not belong to any physical host on the layer-2 domain, then all nodes will silently discard this packet at the layer-3.

Another approach would be to have a management plane command (SNMP or Netconf) to send the magic packet directly to the Ethernet broadcast using any ethertype.

2. Preboot Execution Environment

2.1. Description

Preboot Execution Environment also known as PXE is specified in [PXE21]. It allows for any host to boot a complete viable operating system and file system via the use of Dynamic Host Configuration

Protocol and ancilliary protocols such as Trivial File Transfer Protocol and HyperText Transfer Protocol.

2.2. Issue

The specification has no mention of IPv6 and while DHCP and TFTP support IPv6, there are differences between DHCP for IPv4 and for IPV6. This lack of IPv6 support is addressed in RFC5970 [RFC5970] but there are little to no implementation for this IPv6-enabled PXE. This makes impossible for a thin-client host to boot its complete operating system and file system over an IPv6-only network.

2.3. Mitigation

It is mainly an implementation issue in the boot PROM + DHCPv6 servers. Some of the boot PROMS use flash technology so they could be reprogrammed to fully support RFC5970 [RFC5970]

On the other hand, PXE boot over IPv6 is possible: see [Zimmer2013], relying on Unified Extensible Firmware Interface [UEFI].

3. 3rd issue

Placeholder for any further issue to be described later.

4. IANA Considerations

This document contains no IANA considerations.

5. Security Considerations

The security considerations are detailed in previous sections.

6. Acknowledgements

The author would like to thank Armin Wittman, Alain Fiocco, Steve Simlo, Stig Venaas for some discussions on this topic.

7. References

7.1. Normative References

[PXE21] Intel, , "Preboot Execution Environment (PXE) Specification", 1999, <ftp://download.intel.com/design/archives/wfm/downloads/pxespec.pdf>.

[RFC5970] Huth, T., Freimann, J., Zimmer, V., and D. Thaler, "DHCPv6 Options for Network Boot", RFC 5970, September 2010.

[UEFI] "Unified Extensible Firmware Interface Specification", April 2015, <http://www.uefi.org/sites/default/files/resources/UEFI%202_5.pdf>.

[WOL] AMD, , "Magic Packet Technology", 1995, <http://support.amd.com/TechDocs/20213.pdf>.

[Zimmer2013] Zimmer, V., "Configuring an IPV6 network boot", October 2013, <http://vzimmer.blogspot.com/2013/10/configuring-ipv6-network-boot.html>.

7.2. Informative References

[RFC2644] Senie, D., "Changing the Default for Directed Broadcasts in Routers", BCP 34, RFC 2644, August 1999.

Authors' Addresses

Eric Vyncke
Cisco
De Kleetlaan 6a
Diegem 1831
Belgium

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

Andrew Yourtchenko
Cisco
De Kleetlaan 6a
Diegem 1831
Belgium

Phone: +32478681281
Email: ayourtch@cisco.com

Derk-Jan Valenkamp
ETH Zurich
Weinbergstrasse 43
Zurich 8092
Switzerland

Phone: +41 44 632 50 86
Email: derk-jan.valenkamp@id.ethz.ch