Network Working Group
Request for Comments: 4562
Category: Informational

T. Melsen
S. Blake
Ericsson

June 2006

MAC-Forced Forwarding: A Method for Subscriber Separation on an Ethernet Access Network

Status of This Memo

Copyright Notice

Abstract

This document describes a mechanism to ensure layer-2 separation of Local Area Network (LAN) stations accessing an IPv4 gateway over a bridged Ethernet segment.

The mechanism - called "MAC-Forced Forwarding" - implements an Address Resolution Protocol (ARP) proxy function that prohibits Ethernet Media Access Control (MAC) address resolution between hosts located within the same IPv4 subnet but at different customer premises, and in effect directs all upstream traffic to an IPv4 gateway. The IPv4 gateway provides IP-layer connectivity between these same hosts.

Table of Contents

1. Introduction

The main purpose of an access network is to provide connectivity between customer hosts and service provider access routers (ARs), typically offering reachability to the Internet and other IP networks and/or IP-based applications.

An access network may be decomposed into a subscriber line part and an aggregation network part. The subscriber line - often referred to as "the first mile" - is characterized by an individual physical (or logical, in the case of some wireless technologies) connection to each customer premises. The aggregation network - "the second mile" - performs aggregation and concentration of customer traffic.

The subscriber line and the aggregation network are interconnected by an Access Node (AN). Thus, the AN constitutes the border between individual subscriber lines and the common aggregation network. This is illustrated in the following figure.

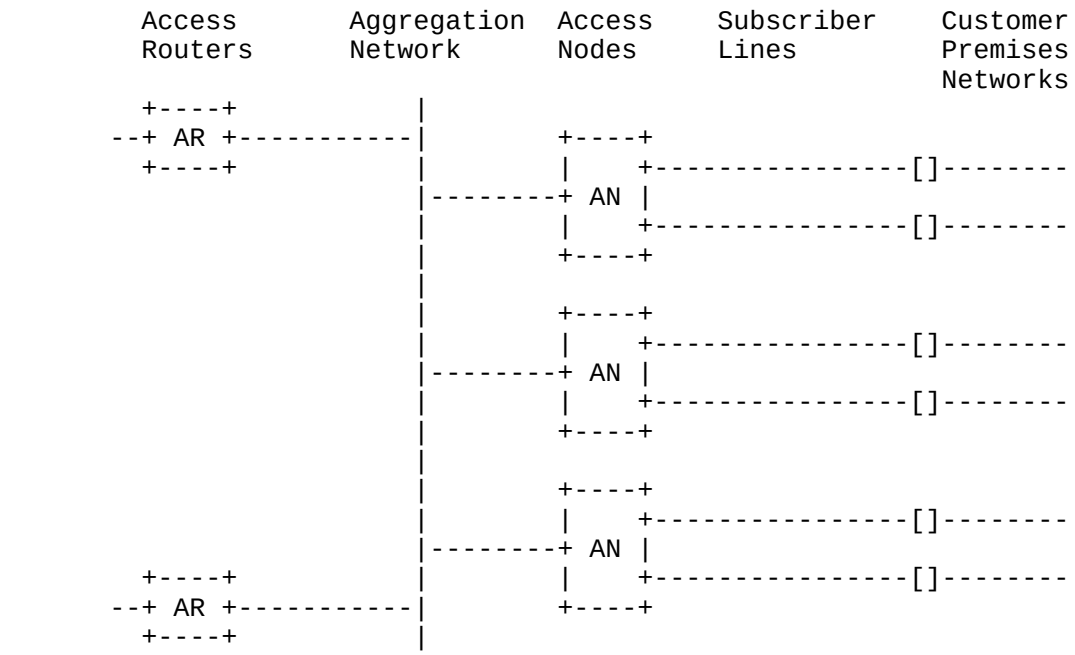

1.1. Access Network Requirements

There are two basic requirements that an access network solution must satisfy:

1. Layer-2 separation between customer premises.

2. High IPv4 address assignment efficiency.

It is required that all traffic to and from customer hosts located at different premises (i.e., accessed via different subscriber lines or via different access networks) be forwarded via an AR, and not bridged or switched at layer-2 (Requirement 1; see also requirement R-40 in [TR101]). This enables the access network service provider to use the AR(s) to perform security filtering, policing, and accounting of all customer traffic. This implies that within the access network, layer-2 traffic paths should not exist that circumvent an AR (with some exceptions; see Section 3.4).

In ATM-based access networks, the separation of individual customer hosts' traffic is an intrinsic feature achieved by the use of ATM permanent virtual connections (PVCs) between the customers' access device (e.g., DSL modem) and the AR (typically co-located/integrated with access control functionality in a Broadband Remote Access Server

(BRAS)). In this case, the AN is an ATM-based Digital Subscriber Line Access Multiplexer (DSLAM).

This document, however, targets Ethernet-based access networks. Techniques other than ATM PVCs must be employed to ensure the desired separation of traffic to and from individual customer hosts.

Efficient address assignment is necessary to minimize consumption of the scarce IPv4 address space (Requirement 2). See [RFC3069] for further discussion. Address assignment efficiency is improved if host addresses are assigned out of one or more large pools, rather than by being assigned out of separate, smaller subnet blocks allocated to each customer premises. IPv6 address assignment efficiency is of much less concern, and it is anticipated that IPv6 deployments will allocate separate IPv6 subnet blocks to each customer premises [v6BB].

1.2. Using Ethernet as an Access Network Technology

A major aspect of using Ethernet as an access technology is that traffic pertaining to different customer hosts is conveyed over a shared broadcast network. Layer-2 isolation between customer premises networks could be provided by implementing access router functionality in each EAN, treating each subscriber line as a separate IP interface. However, there are a variety of reasons why it is often desirable to avoid IP routing in the access network, including the need to satisfy regulatory requirements for direct layer-2 accessibility to multiple IP service providers. In addition, this solution would not solve Requirement 2.

To avoid IP routing within the access network, the Ethernet aggregation network is bridged via EANs to individual Ethernet networks at the customers' premises. If the EANs were standard Ethernet bridges, then there would be direct layer-2 visibility between Ethernet stations (hosts) located at different customers' premises. Specifically, hosts located within the same IP subnet would have this visibility. This violates Requirement 1 (Section 1.1) and introduces security issues, as malicious end-users thereby can attack hosts at other customers' premises directly at the Ethernet layer.

Existing standardized solutions may be deployed to prevent layer-2 visibility between stations:

- PPP over Ethernet [RFC2516]. The use of PPPoE creates individual PPP sessions between hosts and one or more BRASes over a bridged Ethernet topology. Traffic always flows between a BRAS and hosts,

never directly between hosts. The AN can force upstream traffic to flow only to the BRAS initially selected by the host.

- VLAN per-customer premises network [RFC3069]. Traffic to/from each customer premises network can be separated into different VLANs across the aggregation network between the AN and the AR.

Both solutions provide layer-2 isolation between customer hosts, but they are not considered optimal for broadband access networks, because:

- PPPoE does not support efficient multicast: packets must be replicated on each PPPoE session to hosts listening on a specific multicast group. This negates one of the major advantages of using Ethernet (instead of ATM) as an access technology. This is an especially problematic limitation for services such as IPTV, which require high bandwidth per-multicast group (channel), and which may often have hundreds or thousands of listening customer hosts per group.

- Using VLANs to isolate individual customer premises networks also forces multicast packets to be replicated to each VLAN with a listening host. Furthermore, the basic limit of a maximum of 4096 VLANs per-Ethernet network limits the scalability of the solution. This scalability limit can be removed by deploying VLAN stacking techniques within the access network, but this approach increases provisioning complexity.

The solution proposed in this document avoids these problems.

2. Terminology

Access Node (AN)
: The entity interconnecting individual subscriber lines to the shared aggregation network.

Access Router (AR)
: The entity interconnecting the access network to the Internet or other IP-based networks. The AR provides connectivity between hosts on the access network at different customer premises. It is also used to provide security filtering, policing, and accounting of customer traffic.

Application Server (AS)
: A server, usually owned by a service provider, that attaches directly to the aggregation network and is directly reachable at layer-2 by customer hosts.

Ethernet Access Node (EAN)
An Access Node supporting Ethernet-based subscriber lines and uplinks to an Ethernet-based aggregation network and MAC-Forced Forwarding. For example, for xDSL access, the EAN is an Ethernet-centric DSLAM. The EAN is a special type of filtering bridge that does not forward Ethernet broadcast and multicast frames originating on a subscriber line to other subscriber lines, but either discards them or forwards them upstream (towards the aggregation network). The EAN also discards unicast Ethernet frames that originate on a subscriber line and are not addressed to an AR.

3. Solution Aspects

The basic property of the solution is that the EAN ensures that upstream traffic is always sent to a designated AR, even if the IP traffic should ultimately flow between customer hosts located within the same IP subnet.

The solution has three major aspects:

1. Initially, the EAN obtains the IP and MAC addresses of the allowed target ARs for each customer host.

2. The EAN replies to any upstream ARP request [RFC0826] from customer hosts with the MAC address of an allowed target AR.

3. The EAN discards any upstream unicast traffic to MAC addresses other than the allowed target ARs. The EAN also discards all non-essential broadcast and multicast packets received on subscriber lines.

These aspects are discussed in the following sections.

3.1. Obtaining the IP and MAC Addresses of the Access Routers

An access network may contain multiple ARs, and different hosts may be assigned to different (groups of) ARs. This implies that the EAN must register the assigned AR addresses on a per-customer host basis.

For each customer host, one of the ARs is acting as the default gateway. If a customer has simultaneous access to multiple ARs, the other ARs typically will provide access to other IP networks.

The EAN learns the IPv4 address of the allowed target ARs in one of two ways, depending on the host IPv4 address assignment method. For each host using Dynamic Host Configuration Protocol (DHCP), the EAN learns the AR IPv4 addresses dynamically by snooping the DHCPACK

reply to a host [RFC2131]. If a host using DHCP shall have simultaneous access to multiple ARs, DHCP option 121 [RFC3442] or DHCP option 33 [RFC2132] must be used to specify them for that host. If static address assignment is used instead of DHCP, then AR IPv4 addresses must be pre-provisioned in the EAN by the network operator. In both cases, the EAN will ARP to determine the ARs' corresponding MAC addresses. This can be done immediately after the IPv4 addresses are learned or when the MAC addresses are first required.

The DHCP server can associate customer hosts with subscriber lines if the EAN uses the DHCP Relay Agent Information Option (82) to convey a subscriber line identifier to the DHCP server in DHCP messages flowing upstream from the customer host [RFC3046].

3.2. Responding to ARP Requests

If all customer networks were assigned individual IP subnet blocks (and if routing protocols were blocked inside the access network), then all upstream traffic would normally go to an AR (typically the default gateway), and the EAN could validate all upstream traffic by checking that the destination MAC address matched that of an AR.

However, to comply with Requirement 2 of Section 1.1, residential customer networks are not (usually) assigned individual IPv4 subnet blocks. In other words, several hosts located at different premises are within the same IPv4 subnet. Consequently, if a host wishes to communicate with a host at another premises, an ARP request is issued to obtain that host's corresponding MAC address. This request is intercepted by the EAN's ARP proxy, and an ARP reply is sent, specifying an allowed AR MAC address (typically the default gateway's) as the requested layer-2 destination address, in a manner similar to the "proxy ARP" mechanism described in [RFC1812]. In this way, the ARP table of the requesting host will register an AR MAC address as the layer-2 destination for any host within that IPv4 subnet (except those at the same customer premises; see below).

ARP requests for an IPv4 address of an allowed target AR are replied to by the EAN's ARP proxy with that AR's MAC address, rather than the MAC address of the default gateway AR.

An exception is made when a host is ARPing for another host located within the same premises network. If this ARP request reaches the EAN, it should be discarded, because it is assumed to be answered directly by the target host within the premises network. The EAN must keep track of all assigned IPv4 addresses on a subscriber line so that it can detect these ARP requests and discard them.

3.3. Filtering Upstream Traffic

Since the EAN's ARP proxy will always reply with the MAC address of an AR, the requesting host will never learn MAC addresses of hosts located at other premises. However, malicious customers or malfunctioning hosts may still try to send traffic using other unicast destination MAC addresses. The EAN must discard all unicast frames received on a subscriber line that are not addressed to a destination MAC address for an allowed AR (with some exceptions; see Section 3.4.

Similarly, broadcast or multicast packets received on a subscriber line must never be forwarded on other subscriber lines, but only on EAN uplinks to the aggregation network. An EAN must discard all non-ARP broadcast packets received on subscriber lines, except when DHCP is in use, in which case, the EAN must forward client-to-server DHCP broadcast messages (DHCPDISCOVER, DHCPREQUEST, DHCPDECLINE, DHCPINFORM) [RFC2131] upstream. An EAN should rate limit upstream broadcast packets.

Broadcast packets forwarded on an EAN uplink may be forwarded to other EANs by the aggregation network. EANs should discard all broadcast packets received from the aggregation network, except ARPs from ARs for subscriber hosts and server-to-client DHCP messages (DHCPOFFER, DHCPACK, DHCPNAK) [RFC2131], when DHCP is in use.

Filtering of multicast packets to and from an EAN uplink is discussed in Section 6.

3.4. Restricted Access to Application Servers

The previous discussion (Section 3.1) describes how customer hosts are allowed direct layer-2 connectivity only to one or more ARs. Similarly, a customer host could be allowed direct layer-2 access to one or more Application Servers (ASes) which are directly connected to the aggregation network. There is no functional difference in the way MAC-Forced Forwarding treats access to ARs and ASes.

4. Access Router Considerations

Traffic between customer hosts that belong to the same IPv4 subnet but are located at different customer premises will always be forwarded via an AR. In this case, the AR will forward the traffic to the originating network, i.e., on the same interface from where it was received. This normally results in an ICMP redirect message [RFC0792] being sent to the originating host. To prevent this behavior, the ICMP redirect function for aggregation network interfaces must be disabled in the AR.

5. Resiliency Considerations

The operation of MAC-Forced Forwarding does not interfere with or delay IP connectivity recovery in the event of a sustained AR failure. Use of DHCP to configure hosts with information on multiple, redundant ARs, or use of Virtual Router Redundancy Protocol (VRRP) [RFC3768] to implement AR redundancy, allows IP connectivity to be maintained.

MAC-Forced Forwarding is a stateful protocol. If static IPv4 address assignment is used in the access network, then the EAN must be pre-provisioned with state information for the customer hosts which may be reached via a subscriber line, and the ARs associated with those hosts. In the event of a transient EAN failure, the EAN's state database can be quickly recovered from its configuration storage.

If DHCP is used to assign IPv4 addresses in the access network, then MAC-Forced Forwarding operates as a soft-state protocol. Since the DHCP and ARP messages that are snooped to construct the EAN state database are usually sent infrequently, a transient failure may not be detected by either the AR(s) or the customer hosts. Therefore, a transient failure of an EAN could lead to an extended loss of connectivity. To minimize connectivity loss, an EAN should maintain its dynamic state database in resilient storage to permit timely database and connectivity restoration.

The EAN is a single point of attachment between a subscriber line and the aggregation network; hence, the EAN is a single point of connectivity failure. Customers seeking more resilient connectivity should multi-home.

6. Multicast Considerations

Multicast traffic delivery for streams originating within the aggregation network or further upstream and delivered to one or more customer hosts in an access network is supported in a scalable manner by virtue of Ethernet's native multicast capability. Bandwidth efficiency can be enhanced if the EAN behaves as an IGMP snooping bridge; e.g., if it snoops on IGMP Membership Report and Leave Group messages originating on subscriber lines to prune the set of subscriber lines on which to forward particular multicast groups [RFC3376].

An EAN must discard all IPv4 multicast packets received on a subscriber line other than IGMP Membership Report and Leave Group messages [RFC3376]. If a customer host wishes to source multicast packets to a group, the host must tunnel them upstream to a multicast router; e.g., an AR acting as a Protocol Independent Multicast -

Sparse Mode (PIM-SM) Designated Router [RFC2362]. An AR will forward them back into the access network if there are any listening customer hosts.

EAN processing of IPv6 multicast packets is discussed in the next section.

7. IPv6 Considerations

MAC-Forced Forwarding is not directly applicable for IPv6 access networks for the following reasons:

1. IPv6 access networks do not require the same efficiency of address allocation as IPv4 access networks. It is expected that customer premises networks will be allocated unique network prefixes (e.g., /48) accommodating large numbers of customer subnets and hosts [v6BB].

2. IPv6 nodes do not use ARP, but instead use the Neighbor Discovery Protocol [RFC2461] for layer-2 address resolution.

To simultaneously support both IPv6 and MAC-Forced Forwarding for IPv4, an EAN can implement the unicast, broadcast, and multicast filtering rules described in Section 3.3. To correctly perform unicast filtering, the EAN must learn the IPv6 and MAC addresses of the allowed ARs for a particular subscriber line. It can learn these addresses either through static configuration or by snooping Router Discovery messages exchanged between the customer premises router and one or more ARs [RFC2461].

Multicast is an intrinsic part of the IPv6 protocol suite. Therefore, an EAN must not indiscriminately filter IPv6 multicast packets flowing upstream, although it may rate limit them. Detailed IPv6 multicast filtering rules are not discussed in this document.

8. Security Considerations

MAC-Forced Forwarding is, by its nature, a security function, ensuring layer-2 isolation of customer hosts sharing a broadcast access medium. In that sense, it provides security equivalent to alternative PVC-based solutions. Security procedures appropriate for any shared access medium are equally appropriate when MAC-Forced Forwarding is employed. It does not introduce any additional vulnerabilities over those of standard Ethernet bridging.

In addition to layer-2 isolation, an EAN implementing MAC-Forced Forwarding must discard all upstream broadcast packets, except for valid DHCP messages, and ARP requests (which are proxied by the EAN).

In particular, the EAN must discard any DHCP server replies originating on a subscriber line. Further, an EAN may rate limit upstream broadcast DHCP messages.

An EAN implementing MAC-Forced Forwarding must keep track of IPv4 addresses allocated on subscriber lines. Therefore, the EAN has sufficient information to discard upstream traffic with spoofed IPv4 source addresses.

9. Acknowledgements

The authors would like to thank Ulf Jonsson, Thomas Narten, James Carlson, Rolf Engstrand, Tomas Thyni, and Johan Kolhi for their helpful comments.

10. References

10.1. Normative References

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.

[RFC0826] Plummer, D., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware", STD 37, RFC 826, November 1982.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.

[RFC2362] Estrin, D., Farinacci, D., Helmy, A., Thaler, D., Deering, S., Handley, M., Jacobson, V., Liu, C., Sharma, P., and L. Wei, "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification", RFC 2362, June 1998.

[RFC3046] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[RFC3442] Lemon, T., Cheshire, S., and B. Volz, "The Classless Static Route Option for Dynamic Host Configuration Protocol (DHCP) version 4", RFC 3442, December 2002.

10.2. Informative References

[RFC1812] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.

[RFC3768] Hinden, R., "Virtual Router Redundancy Protocol (VRRP)", RFC 3768, April 2004.

[RFC2461] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

[RFC2516] Mamakos, L., Lidl, K., Evarts, J., Carrel, D., Simone, D., and R. Wheeler, "A Method for Transmitting PPP Over Ethernet (PPPoE)", RFC 2516, February 1999.

[RFC3069] McPherson, D. and B. Dykes, "VLAN Aggregation for Efficient IP Address Allocation", RFC 3069, February 2001.

[TR101] DSL Forum, "Migration to Ethernet-Based DSL Aggregation", Technical Report TR-101, April 2006.

[v6BB] Asadullah, S., Ahmed, A., Popoviciu, C., Savola, P., and J. Palet, "ISP IPv6 Deployment Scenarios in Broadband Access Networks", Work in Progress.

Authors' Addresses

Torben Melsen
Ericsson
Faelledvej
Struer DK-7600
Denmark

EMail: Torben.Melsen@ericsson.com

Steven Blake
Ericsson
920 Main Campus Drive
Suite 500
Raleigh, NC 27606
USA

Phone: +1 919 472 9913
EMail: steven.blake@ericsson.com

Full Copyright Statement

Intellectual Property

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).