

Network Working Group
INTERNET-DRAFT
Category: Best Current Practice
<[draft-aboba-dhc-nad-ipv4-00.txt](#)>
[12](#) June 2003

Bernard Aboba
Microsoft Corporation

IPv4 Network Attachment Detection

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC 2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

This specification attempts to synthesize experience garnered over the years in the deployment of hosts supporting ARP, DHCP and IPv4 Link-Local addresses. Given this experience, this document suggests optimizations for IPv4 network attachment detection as well as heuristics for determining when assignment of an IPv4 Link-Local address is appropriate.

[1](#). Introduction

This draft attempts to synthesize experience garnered over the years in the deployment of hosts supporting ARP [[RFC826](#)], DHCP [[RFC2131](#)], and Link-Local IPv4 addresses [[IPv4LL](#)]. There are several reasons which this experience is valuable:

INTERNET-DRAFT

IPv4 NAD

12 June 2003

- [a] Avoiding inappropriate use of Link-Local IPv4 addresses. Experience has shown that in the vast majority of cases, the assignment of Link-Local IPv4 addresses is inappropriate. That is, the IPv4 host assigning an Link-Local IPv4 address either is not connected to a network at all, in which case assignment of an Link-Local IPv4 address does no good; or the host is in fact present on a network with a DHCPv4 server but for one reason or another does not receive a response to a DHCPREQUEST or DHCPDISCOVER.
- [b] Optimization of network attachment detection. The time required to detect movement (or lack of movement) between subnets, and to obtain (or continue to use) a valid IPv4 address represents a significant fraction of the overall latency resulting from movement between points of attachment on the network. As a result, optimization of network attachment detection in IPv4 hosts is helpful, to the extent that it is achievable.

In order to provide suggestions for handling problems [a] and [b], this document suggests the following basic principles:

- [1] Utilization of link layer hints. Link layers such as IEEE 802 [IEEE802] provide hints about whether a host remains on the same subnet despite changing its point of attachment, or even whether the host is connected to an adhoc or infrastructure network. Where available, these hints can be used to guide host behavior - with the understanding that they are not infallible and therefore that the host should be capable of making the correct determination even in the presence of misleading hints. Link layer hints are described in more detail in [Section 3](#).
- [2] Link-Local IPv4 addressing as a mechanism of last resort. According to [IPv4LL], once a Link-Local IPv4 address is assigned, the DHCPv4 server may not be queried again for 5 minutes. As a result, the inappropriate assignment of a Link-Local IPv4 address results in an extended period of limited connectivity. For a host that may change its point of attachment more frequently than every 5 minutes, the inappropriate assignment of an Link-Local IPv4 address is more than just an annoyance - it can result in an ongoing inability to connect. As a result, this document suggests that hosts behave conservatively with respect to assignment of Link-Local IPv4 addresses, using them only when there is good evidence that a DHCPv4 server is not present.

1.1. Framework and assumptions

This document specifies a procedure to be performed for IPv4 network attachment detection that depends on two phases: a connectivity test phase, and an IPv4 address acquisition phase. The tests specified in

this document are conducted on connection to a network. On disconnection from a network, there is no need to take action until the host is reconnected, since it is typically not possible for a host to communicate until it has obtained connectivity. Therefore, contrary to [\[RFC2131\] Section 3.7](#), no profitable actions can be taken on network disconnection.

The purpose of the connectivity test phase is for the host to be able to quickly determine whether it remains connected to a network on which it had previously obtained a still valid routable IPv4 address. In this case it is possible to continue to use the existing IPv4 address without having to reacquire it.

If the connectivity test is not conclusive, then the host may have moved subnets, so that it attempts to reacquire its IPv4 address. It is possible that the attempt to reacquire an address will be unsuccessful, either due to lack of receipt of a DHCPACK, or because a DHCPNAK is received. In this case, the host attempts to acquire an IPv4 address. If the host does not have a valid routable IPv4 address on the "most likely" network, then the connectivity test is skipped and the host attempts to acquire an IPv4 address.

1.2. Connectivity Test

Where the host has not yet confirmed whether it has moved subnets, the first step is to test connectivity to the default gateway on the "most likely point of attachment".

Where the host has retained the IP and MAC address of the default gateway on the "most likely" network, an attempt is made to demonstrate connectivity to the primary default gateway. If connectivity cannot be demonstrated, the host proceeds to IPv4 address acquisition (or re-acquisition), rather than testing connectivity to alternate gateways.

In order to determine the "most likely point of attachment" it is assumed that the host is capable of obtaining and writing to stable

storage parameters relating to networks that it connects to. This information includes the following:

- [1] IP and MAC address(es) of the default gateway(s).
- [2] Link layer information associated with each network. Link layer information is discussed in more detail in [Section 2](#).

Based on this information, the host may be able to make an educated guess as to whether it is likely to have moved between subnets, and if so, to which network it has moved. Link layer information which can be used to determine the "most likely point of attachment" is discussed in

[Section 3](#). Where no additional information is available, by default the host assumes that the "most likely point of attachment" is the network to which it was most recently connected.

Demonstrating connectivity to the default gateway on the "most likely point of attachment" is useful only where the host can immediately utilize a valid routable IPv4 address where connectivity to be confirmed. If the host does not have a valid routable IPv4 address on the "most likely" network, then it will need to obtain an IPv4 address, so that it is preferable to skip the connectivity test entirely and attempt IPv4 address acquisition or reacquisition.

Where the host had previously obtained a Link-Local IPv4 address on the "most likely" network it is also preferable to skip the connectivity test since Link-Local IPv4 addresses are only used as a last resort. Therefore IPv4 address acquisition or reacquisition is preferred in this case as well.

Where the host has evidence that it has moved subnets (such as association to a different SSID), but no valid routable IPv4 address on the new network or information on the default gateway on that network, the connectivity test is also skipped since insufficient information is available to conduct it.

To perform the connectivity test, an ARP Request SHOULD be sent, using the host's MAC address as the source, and the MAC address of the primary default gateway as the destination. The host sets the target protocol address (ar\$tpa) to the IPv4 address of the primary default gateway, and uses its own MAC address in the sender hardware address field (ar\$sha).

Since the host does not yet know whether it has moved subnets, it MUST set the sender protocol address field (ar\$spa) to 0.0.0.0. This prevents poisoning of the ARP cache with a (potentially invalid) sender IPv4 address.

The ARP Request is sent to the unicast MAC address of the default gateway rather than to the broadcast address in order to confirm both the IPv4 address and MAC address of the default gateway. This allows the host to confirm connectivity to the default gateway even where the host moves between two private networks, since in this case the IPv4 address of the default gateway could remain the same, while the MAC address would change.

Sending an ICMP Echo Request to the default gateway IPv4 address does not provide the same level of assurance and SHOULD NOT be substituted for an ARP Request sent to the MAC address of the default gateway. Where a host moves from one private network to another, an ICMP Echo Request can result in an ICMP Echo Response even when the default gateway has changed, as long as the IPv4 address remains the same. This

can occur, for example, where a host moves from one home network using prefix 192.168/16 to another one. In addition, if the ping is sent with TTL > 1, then an ICMP Echo Response can be received from an offlink gateway.

If a valid ARP Response is received, the MAC address in the target hardware address field (ar\$tha) and the IPv4 address in the target protocol address field (ar\$tpa) are matched against the list of networks and associated default gateway parameters. If a match is found, then if the host has a valid IPv4 address lease on the matched network, the host continues to use that IPv4 address, subject to the lease reacquisition and expiration behavior described in [\[RFC2131\], Section 4.4.5](#).

If the initial ARP Request does not elicit a Response, the host waits 1 second and then sends an ARP Request to the broadcast address. If no ARP Response is received in response to this second (broadcast) Request, the host proceeds to the next phase. If a valid ARP Response is received, but cannot be matched against known networks, the host assumes it has moved subnets and moves on to the next phase.

[1.3](#). IP address acquisition

Where the connectivity test is inconclusive, where the host had previously obtained an IPv4 Link-Local address on the "most likely" network, or where the host has no valid IPv4 address lease on the "most likely" network, the host attempts to obtain an IPv4 address.

If the host has a valid cached configuration but is unable to confirm connectivity to default gateway on the "most likely point of attachment" then the host seeks to verify the cached configuration by entering the INIT-REBOOT state, and sending a DHCPREQUEST to the broadcast address as specified in [\[RFC2131\] Section 4.4.2](#).

If the host does not have a valid cached configuration, or it had previously obtained a Link-Local IPv4 address on the "most likely" network, then the host enters the INIT state and sends a DHCPDISCOVER packet to the broadcast address, as described in [\[RFC2131\] Section 4.4.1](#). If the host does not receive a response to a DHCPREQUEST or DHCPDISCOVER, then it retransmits as specified in [\[RFC2131\] Section 4.1](#).

As discussed in [\[RFC2131\], Section 4.4.4](#), a host in INIT or REBOOTING state that knows the address of a DHCP server may use that address in the DHCPDISCOVER or DHCPREQUEST rather than the IP broadcast address. However, sending a DHCPREQUEST to the unicast address when in INIT-REBOOT state is not appropriate since it is possible that the client has moved to another subnet, and therefore the DHCPREQUEST needs to be forwarded to and from the DHCP server by a DHCP Relay so that the

response can be broadcast. This ensures that the host will receive a response regardless of whether the cached IP address is correct for the network to which it has connected.

As noted in [\[RFC2131\] Section 3.2](#), if the host possesses a valid routable IPv4 address on the "most likely" network and does not receive a response after employing the retransmission algorithm, the client MAY choose to use the previously allocated network address and configuration parameters for the remainder of the unexpired lease. This is preferable to assigning a Link-Local IPv4 address if the host has reason to believe that it is connected to a network on which it possesses a valid IPv4 address lease. This would be the case, for example, where a host reconnects to an IEEE 802.11 network with the same SSID as a network on which it had previously obtained a still-valid IPv4 address lease.

Alternatively, if the host does not have a valid IPv4 address lease on

the "most likely" network and does not receive a response after employing the retransmission algorithm, it MAY assign a Link-Local IPv4 address. This would be the preferred behavior, for example, in situations where the host connects to an adhoc IEEE 802.11 network, unless a routable IPv4 address had previously been assigned on that network.

However even in this situation, it is likely that the failure to obtain a routable IPv4 address represents a temporary aberration, rather than legitimate detection of an adhoc network. In such a circumstance, it is therefore desirable to abandon the assignment of an Link-Local IPv4 address as soon as a valid IPv4 address lease can be obtained. As a result, it is RECOMMENDED that in such a case, the host will attempt to obtain an IPv4 address assignment at 30 second intervals.

1.4. Link layer hints

In order to assist in IPv4 network attachment detection, link layer information associated with each network may be retained by the host. Based on information obtained from the link layer, the host may be able to make an educated guess as to whether it has moved between subnets, or remained on the same subnet. If it is likely to have moved between subnets, the host may have an educated guess as to which subnet it has moved to.

For networks running over PPP [[RFC1661](#)], this can include the link characteristics negotiated in LCP, the IP parameters negotiated in IPCP, and perhaps the associated phone number.

On IEEE 802 wired networks, source of link layer information include link-layer discovery traffic as well as information exchanged as part of IEEE 802.1X authentication. Link-layer discovery traffic includes Cisco

CDP exchanges as well as network identification information passed in the EAP-Request/Identity or within an EAP method exchange. IEEE 802.1 is currently working on standardization of Link Layer Discovery Protocol (LLDP) [LLDP] so that this may also be taken into account.

For example, Cisco CDP advertisements can provide information on the IP address of the device, allowing for an intelligent guess as to the likely subnet to which it is connected. When used with IEEE 802.1X authentication, the EAP-Request/Identity exchange may contain the name

of the authenticator, also providing information on the potential network. Similarly, during the EAP method exchange the authenticator may supply information that may be helpful in identifying the network to which the device is attached.

In IEEE 802.11, [[IEEE80211](#)] the device provides information in Beacon and/or Probe Response messages, such as the SSID, BSSID, and capabilities. It also includes information on whether the network is operating in Infrastructure or adhoc mode. As described in [Congdon], it is possible to assign a Station to a VLAN dynamically, based on the results of IEEE 802.1X [[IEEE8021X](#)] authentication. This implies that a single SSID may offer access to multiple VLANs, and in practice most large WLAN deployments offer access to multiple subnets. Thus, associating to the same SSID is a necessary, but not necessarily a sufficient condition for remaining within the same subnet. In order to provide additional guidance on the subnets to which a given AP offers access, additional subnet-related Information Elements (IEs) have been proposed for addition to the IEEE 802.11 Beacon and Probe Response messages [handoffIE].

While a Station associating to the same SSID may not necessarily remain within the same subnet, on the other hand, a Station associated to a different SSID is likely to have changed subnets. As a result, a Station associating with a different SSID MAY forgo the "remains connected" step in Section ? and go straight to the "Discover new address" step.

[2.](#) Normative References

- [RFC791] Postel, J., "Internet Protocol", [RFC 791](#), USC/Information Sciences Institute, September 1981.
- [RFC792] Postel, J., "Internet Control Message Protocol", [RFC 792](#), USC/Information Sciences Institute, September 1981.
- [RFC1256] Deering, S., "ICMP Router Discovery Messages", [RFC 1256](#), Xerox PARC, September 1991.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March, 1997.

- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), March 1997.
- [RFC2132] Alexander, S. and Droms, R., "DHCP Options and BOOTP Vendor Extensions", [RFC 2132](#), Silicon Graphics, Inc., Bucknell University, March 1997.
- [IPv4LL] Cheshire, S., Aboba, B. and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", Internet draft (work in progress), [draft-ietf-zeroconf-ipv4-linklocal-08.txt](#), June 2003.

3. Informational References

- [RFC1034] Mockapetris, P., "Domain Names - Concepts and Facilities", STD 13, [RFC 1034](#), USC/Information Sciences Institute, November 1987.
- [RFC1035] Mockapetris, P., "Domain Names - Implementation and Specification", STD 13, [RFC 1035](#), USC/Information Sciences Institute, November 1987.
- [RFC1058] Hedrick, C.L., "Routing Information Protocol", [RFC 1058](#), Rutgers University, June 1, 1988.
- [RFC1332] McGregor, G., "PPP Internet Control Protocol", [RFC 1332](#), Merit, May 1992.
- [RFC1661] Simpson, W., Editor, "The Point-to-Point Protocol (PPP)", STD 51, [RFC 1661](#), Daydreamer, July 1994.
- [RFC1877] Cobb, S., "PPP Internet Protocol Control Protocol Extensions for Name Server Addresses", [RFC 1877](#), December 1995.
- [RFC2284] Blunk, L. and J. Vollbrecht, "PPP Extensible Authentication Protocol (EAP)", [RFC 2284](#), March 1998.
- [RFC3118] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", [RFC 3118](#), June 2001.
- [IEEE8021X] IEEE Standards for Local and Metropolitan Area Networks: Port based Network Access Control, IEEE Std 802.1X-2001, June 2001.

- [RFC2434] Alvestrand, H. and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 2434](#), October 1998.
- [IEEE802] IEEE Standards for Local and Metropolitan Area Networks: Overview and Architecture, ANSI/IEEE Std 802, 1990.
- [IEEE8021Q] IEEE Standards for Local and Metropolitan Area Networks: Draft Standard for Virtual Bridged Local Area Networks, P802.1Q, January 1998.
- [IEEE8023] ISO/IEC 8802-3 Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Common specifications - Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications, (also ANSI/IEEE Std 802.3-1996), 1996.
- [IEEE80211] Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE Std. 802.11-1999, 1999.

Acknowledgments

The authors would like to acknowledge Erik Guttman and Erik Nordmark of Sun Microsystems, Ted Lemon of Nominum and Thomas Narten of IBM for contributions to this document.

Authors' Addresses

Bernard Aboba
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

E-Mail: bernarda@microsoft.com
Phone: +1 425 706 6605
Fax: +1 425 936 7329

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain

to the implementation or use of the technology described in this

INTERNET-DRAFT

IPv4 NAD

12 June 2003

document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved. This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English. The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns. This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Expiration Date

This memo is filed as <[draft-aboba-dhc-nad-ipv4-00.txt](#)>, and expires December 22, 2003.

Aboba

Informational

[Page 10]