Network Working Group
Request for Comments: 4957
Category: Informational

S. Krishnan, Ed.
Ericsson Research
N. Montavont
GET ENST Bretagne
E. Njedjou
France Telecom
S. Veerepalli
Qualcomm
A. Yegin, Ed.
Samsung

August 2007

Link-Layer Event Notifications for Detecting Network Attachments

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

Certain network access technologies are capable of providing various types of link-layer status information to IP. Link-layer event notifications can help IP expeditiously detect configuration changes. This document provides a non-exhaustive catalogue of information available from well-known access technologies.

Table of Contents

1. Introduction

It is not an uncommon occurrence for a node to change its point of attachment to the network. This can happen due to mobile usage (e.g., a mobile phone moving among base stations) or nomadic usage (e.g., road-warrior case).

A node changing its point of attachment to the network may end up changing its IP subnet and therefore require reconfiguration of IP-layer parameters, such as IP address, default gateway information, and DNS server address. Detecting the subnet change can usually use network-layer indications (such as a change in the advertised prefixes for IPv6). But such indications may not be always available (e.g., Detecting Network Attachment in IPv6 (DNAv6)) to the node upon changing its point of attachment.

Link-layer event notifications can help IP expeditiously detect configuration changes. This document provides a non-exhaustive catalog of information available from some access technologies, and discusses the interpretation of this information at the IP layer. This document is not intended to specify or change the behavior of these access technologies in any manner.

Additional information can be conveyed along with the event, such as the identifier of the network attachment point (e.g., IEEE 802.11 Basic Service Set Identification (BSSID) and Service Set Identifier (SSID)), or network-layer configuration parameters obtained via the link-layer attachment process if available. It is envisaged that such event notifications can in certain circumstances be used to expedite the inter-subnet movement detection and reconfiguration process. For example, the notification indicating that the node has established a new link-layer connection may be used for immediately probing the network for a possible configuration change. In the absence of such a notification from the link layer, IP has to wait for indications that are not immediately available, such as receipt of the next scheduled router advertisement, unreachability of the default gateway, etc.

It should be noted that a link-layer event notification does not always translate into a subnet change. Even if the node has torn down a link-layer connection with one attachment point and established a new connection with another, it may still be attached to the same IP subnet. For example, several IEEE 802.11 access points can be attached to the same IP subnet. Moving among these access points does not warrant any IP-layer configuration change.

In order to enable an enhanced scheme for detecting change of subnet, we need to define link-layer event notifications that can be realistically expected from various access technologies. The objective of this document is to provide a catalogue of link-layer events and notifications in various architectures. While this document mentions the utility of this information for detecting change of subnet (or, detecting network attachment - DNA), the detailed usage is left to other documents, namely, DNA solution specifications.

The document limits itself to the minimum set of information that is necessary for solving the DNA problem [RFC4135]. A broader set of information (e.g., signal strength, packet loss, etc.) and events (e.g. link down) may be used for other problem spaces, such as anticipation-based Mobile IP fast handovers [RFC4881], [RFC4068], etc.

These event notifications are considered with hosts in mind, although they may also be available on the network side (e.g., on the access points and routers). An API or protocol-based standard interface may be defined between the link layer and IP for conveying this information. That activity is beyond the scope of this document.

2. Terminology

Link: is a communication facility or medium over which network nodes can communicate. Each link is associated with a minimum of two endpoints. An "attachment point" is the link endpoint on the link to which the node is currently connected, such as an access point, a base station, or a wired switch.

Link up: is an event provided by the link layer that signifies a state change associated with the interface becoming capable of communicating data packets. This event is associated with a link-layer connection between the node and an attachment point.

BSSID: Basic Service Set Identification

DNA: Detecting Network Attachment

GPRS: General Packet Radio Service

PDP: Packet Data Protocol

SSID: Service Set Identifier

3. Link-Layer Event Notifications

Link-layer event notifications are considered to be one of the inputs to the DNA process. A DNA process is likely to take other inputs (e.g., presence of advertised prefixes, reachability of default gateways) before determining whether IP-layer configuration must be updated. It is expected that the DNA process can take advantage of link-layer notifications when they are made available to IP. While by itself a link-layer notification may not constitute all the input DNA needs, it can at least be useful for prompting the DNA process to collect further information (i.e., other inputs to the process). For example, the node may send a router solicitation as soon as it learns that a new link-layer connection is established.

The link-layer event that is considered most useful to DNA process is the link up event. The associated notifications can be provided to the IP-layer after the event concludes successfully. The link up events and notifications are associated with a network interface on the node. The IP module may receive simultaneous independent notifications from each one of the network interfaces on the node.

The actual event is managed by the link layer of the node through execution of link-layer protocols and mechanisms. Once the event successfully completes within the link layer, its notification is delivered to the IP-layer. By the time the notification is delivered, the link layer of the node must be ready to accept IP packets from the IP and the physical layers. Each time an interface changes its point of attachment, a link up event should be generated.

There is a non-deterministic usage of the link up notification to accommodate implementations that desire to indicate the link is up, but the data transmission may be blocked in the network (see IEEE 802.3 discussion). A link up notification may be generated with an appropriate attribute, conveying its non-deterministic nature, to convey the event. Alternatively, the link-layer implementation may choose to delay the link up notification until the risk conditions cease to exist.

If a non-deterministic link up was generated, another link up must follow as soon as the link layer is capable of generating a deterministic notification. The event attributes may indicate whether the packets transmitted since the previous notification were presumed to be blocked or allowed by the network, if the link layer could determine the exact conditions.

The deterministic link up event following a non-deterministic link up event can be treated differently by consumers of the link up event. For example, the second link up event need not trigger a confirmation process, if the first one already did.

A node may have to change its IP-layer configuration even when the link-layer connection stays the same. An example scenario is the IPv6 subnet renumbering [RFC2461]. Therefore, there exist cases where IP-layer configuration may have to change even without the IP layer receiving a link up notification. Therefore, a link-layer notification is not a mandatory indication of a subnet change.

A link up notification may optionally deliver information relating to the attachment point. Such auxiliary information may include the identity of the attachment point (e.g., base station identifier), or the IP-layer configuration parameters associated with the attached subnet (e.g., subnet prefix, default gateway address, etc.). While merely knowing that a new link-layer connection is established may prompt the DNA process to immediately seek other clues for detecting a network configuration change, auxiliary information may constitute further clues (and even the final answers sometimes). In cases where there is a one-to-one mapping between the attachment point identifiers and the IP-layer configurations, learning the former can reveal the latter. Furthermore, IP-layer configuration parameters obtained during the link-layer connection may be exactly what the DNA process is trying to discover.

The link-layer process leading to a link up event depend on the link technology. While a link-layer notification must always indicate that the link up event occurred, the availability and types of auxiliary information on the attachment point depends on the link-layer technology as well. The following subsections examine four link-layer technologies and describe when a link-layer notification is generated and what information is included in it.

3.1. GPRS/3GPP

GSM Packet Radio System (GPRS) provides packet-switched data transmission over a cellular network [GPRS][GPRS-LINK].

The GPRS architecture consists of a Radio Access Network and a packet domain Core Network.

- The GPRS Radio Access Network is composed of Mobile Terminals (MTs), a Base Station Subsystem and Serving GPRS Support Nodes (SGSNs).

- An IP Core Network that acts as the transport backbone of user datagrams between SGSNs and Gateway GPRS Support Nodes (GGSNs). The GGSN ensures the GPRS IP core network connectivity with external networks, such as the Internet or Local Area Networks. The GGSN acts as the default IP gateway for the MT.

A GPRS MT that wants to establish IP connectivity establishes first a connection to the GPRS network and one or more PDP Context associations between the MT and the GGSN. It is only after the PDP Context has been established and after address autoconfiguration and tunneling mechanism have taken place that the MT's IP packets can be forwarded to and from its remote IP peers. The aim of PDP Context establishment is also to provide IP-level configuration on top of the GPRS link-layer attachment.

Successful establishment of a PDP Context on a GPRS link signifies the availability of IP service to the MT. Therefore, this link-layer event generates a link up event notification sent to the IP layer.

An MT may establish a secondary PDP Context while reusing the IP configuration acquired from a previously established and active PDP Context. Such a secondary PDP Context does not provide additional information to the IP layer and only allows another quality-of-service (QoS) profile to be used. The activation of such a secondary PDP context does not usually generate a link up event since it does not require new IP parameters. However, other additional PDP Context activations are to be treated as indicated earlier.

With IPv4, the auxiliary information carried along with this notification is the IPv4 address of the MT that is obtained as part of the PDP Context. With IPv6, the PDP Context activation response does not come along with a usable IPv6 address. Effectively, the IPv6 address received from the GGSN in the PDP address field of the message does not contain a valid prefix. The MN actually only uses the interface identifier extracted from that field to form a link-local address that it uses afterwards to obtain a valid prefix (e.g., by stateless [RFC2462][GPRS-CN] or stateful [RFC3315] [GPRS-GSSA] address configuration). Therefore, no IPv6-related auxiliary information is provided to the IP layer.

3.2. cdma2000/3GPP2

cdma2000-based 3GPP2 packet data services provide mobile users wide area high-speed access to packet switched networks [CDMA2K]. Some of the major components of the 3GPP2 packet network architecture consist of:

- Mobile Station (MS), which allows mobile access to packet-switched networks over a wireless connection.

- Radio Access Network, which consists of the Base Station Transceivers, Base Station Controllers, and the Packet Control Function.

- Network Access Server known as the Packet Data Switching Node (PDSN). The PDSN also serves as default IP gateway for the IP MS.

3GPP2 networks use the Point-to-Point Protocol (PPP [RFC1661]) as the link-layer protocol between the MS and the PDSN. Before any IP packets may be sent or received, PPP must reach the Network-Layer Protocol phase, and the IP Control Protocol (IPCP [RFC1332], IPV6CP [RFC2472]) must reach the Opened state. When these states are reached in PPP, a link up event notification is delivered to the IP layer.

When the PPP is used for 3GPP2 Simple (i.e., non-Mobile) IPv4 Service, IPCP enables configuration of an IPv4 address on the MS. This IPv4 address is provided as the auxiliary information along with the link up notification. IPV6CP used for Simple IPv6 service does not provide an IPv6 address, but the interface identifiers for local and remote endpoints of the PPP link. Since there is no standards-mandated correlation between the interface identifier and other IP-layer configuration parameters, this information is deemed not useful for DNA (nevertheless, it may be provided as auxiliary information for other uses).

3.3. IEEE 802.11/WiFi

IEEE 802.11-based WiFi networks are the wireless extension of the Local Area Networks. Currently available standards are IEEE 802.11b [IEEE-802.11b], IEEE 802.11g [IEEE-802.11g], and IEEE 802.11a [IEEE-802.11a]. The specifications define both the MAC layer and the physical layer. The MAC layer is the same for all these technologies.

Two operating modes are available in the IEEE 802.11 series, either infrastructure mode or ad-hoc mode. In infrastructure mode, all link-layer frames are transmitted to an access point (AP) that then forwards them to the final receiver. A station (STA) establishes an IEEE 802.11 association with an AP in order to send and receive IP packets. In a WiFi network that uses Robust Secure Network (RSN [IEEE-802.11i]), successful completion of the 4-way handshake between the STA and AP commences the availability of IP service. The link up

event notification is generated upon this event. In non-RSN-based networks, successful association or re-association events on the link layer causes a link up notification sent to the IP layer.

As part of the link establishment, the STA learns the BSSID and SSID associated with the AP. The BSSID is a unique identifier of the AP, usually set to the MAC address of the wireless interface of the AP. The SSID carries the identifier of the Extended Service Set (ESS) -- the set composed of APs and associated STAs that share a common distribution system. The BSSID and SSID may be provided as auxiliary information along with the link up notification. Unfortunately, this information does not provide a deterministic indication of whether the IP-layer configuration must be changed upon movement. There is no standards-mandated one-to-one relation between the BSSID/SSID pairs and IP subnets. An AP with a given BSSID can connect a STA to any one of multiple IP subnets. Similarly, an ESS with the given SSID may span multiple IP subnets. And finally, the SSIDs are not globally unique. The same SSID may be used by multiple independent ESSs. Nevertheless, BSSID/SSID information may be used in a probabilistic way by the DNA process; hence, it is provided with the link up event notification.

In ad-hoc mode, mobile stations (STA) in range may directly communicate with each other, i.e., without any infrastructure or intermediate hop. The set of communicating STAs is called IBSS for Independent Basic Service Set. In an IBSS, only STA services are available, i.e., authentication, deauthentication, privacy, and MAC Service Data Unit (MSDU) delivery. STAs do not associate with each other, and therefore may exchange data frames in state 2 (authenticated and not associated) or even in state 1 (unauthenticated and unassociated) if the Distribution System is not used (i.e., "To DS" and "From DS" bits are clear). If authentication is performed, a link up indication can be generated upon authentication. Concerning the link layer identification, both the BSSID (which is a random MAC address chosen by a STA of the IBSS) and SSID may be used to identify a link, but not to make any assumptions on the IP network configuration.

3.4. IEEE 802.3 CSMA/CD

IEEE 802.3 CSMA/CD (commonly referred to as Ethernet) is the most commonly deployed Local Area Network technology in use today. As deployed today, it is specified by a physical layer/medium access control (MAC) layer specification [IEEE-802.3]. In order to provide connection of different LANs together into a larger network, 802.3 LANs are often bridged together [IEEE-802.1D].

In this section, the terms 802.3 and Ethernet are used interchangeably. This section describes some issues in providing link-layer indications on Ethernet networks, and shows how bridging affects these indications.

In Ethernet networks, hosts are connected by wires or by optic fibre to a switch (bridge), a bus (e.g., coaxial cable), a repeater (hub), or directly to another Ethernet device. Interfaces are symmetric, in that while many different physical layers may be present, medium access control is uniform for all devices.

In order to determine whether the physical medium is ready for frame transfer, IEEE 802.3 Ethernet specifies its own link monitoring mechanism, which is defined for some, but not all, classes of media. Where available, this Link Integrity Test operation is used to identify when packets are able to be received on an Ethernet segment. It is applicable to both wired and optical physical layers, although details vary between technologies (link pulses in twisted pair copper, light levels in fibre).

3.4.1. Link Integrity Tests in 802.3 Networks

Link Integrity Tests in 802.3 networks typically occur at initial physical connection time (for example, at the auto-negotiation stage) and periodically afterwards. They make use of physical-layer specific operations to determine if a medium is able to support link-layer frames [IEEE-802.3].

The status of the link as determined by the Link Integrity Test is stored in the variable 'link_status'. Changes to the value of link_status (for example due to Link Integrity Test failure) will generate link indications if the technology-dependent interface is implemented on an Ethernet device [IEEE-802.3].

The link_status has possible values of FAIL, READY, and OK. In FAIL state, Link Integrity Tests have failed. In READY state, the link segment has passed integrity tests, but auto-negotiation has not completed. In OK state, the medium is able to send and receive packets.

Upon transition to a particular state, the Physical Medium Attachment subsystems generates a PMA_LINK.indicate(link_status). Indications of OK state may be used to generate a link up event notification. These indications do not definitively ensure that packets will be able to be received through the bridge domain, though (see the next section). Such operations are governed by bridging.

3.4.2. IEEE 802.1D Bridging and Its Effects on Link-layer Event Notifications

Ethernet networks commonly consist of LANs joined together by transparent bridges (usually implemented as switches). Transparent bridges require the active topology to be loop free. This is achieved through the Spanning Tree Protocol (STP) or the Rapid Spanning Tree Protocol (RSTP). These protocols exchange Bridge Protocol Data Units (BPDUs), as defined in [IEEE-802.1D]; this leads to the blocking of ports (i.e., not forwarding), where required.

By default, the spanning tree protocol does not know whether a particular newly connected piece of Ethernet will cause a loop.

Therefore, it will block all traffic from and to newly connected ports with the exception of some unbridged management frames. The STP will determine if the port can be connected to the network in a loop-free manner.

For these technologies, even though the link layer appears available, no data packet forwarding will occur until it is determined that the port can be connected to the network in a loop-free environment.

For hosts that are providing indications to upper-layer protocols, even if the host itself does not implement bridging or STP, packet delivery across the network can be affected by the presence of bridges.

A host connected to a bridge port does not receive any explicit indication that the bridge has started forwarding packets. Therefore, a host may not know when STP operations have completed, or when it is safe to inform upper layers to transmit packets.

Where it is not known that forwarding operations are available, a host should assume that RSTP or STP is being performed. Hosts may listen to STP/RSTP and 802.1AB messages to gain further information about the timing of full connectivity on the link, for example, to override an existing indication.

Notably, though, it is not easy for a host to distinguish between disabled bridge ports and non-bridge ports with no active transmitters on them, as Disabled ports will have no traffic on them, and incur 100% sender loss.

If no bridge configuration messages are received within the Bridge_Max_Age interval (default 20s) then it is likely that there is no visible bridge whose port is enabled for bridging (S8.4.5 of [IEEE-802.1D]), since at least two BPDU hello messages would have

been lost. Upon this timeout, a link up notification is generated, if one has not been already.

If a BPDU is received, and the adjacent bridge is running the original Spanning Tree Protocol, then a host cannot successfully send packets until at least twice the ForwardDelay value in the received BPDU has elapsed. After this time, a link up notification is generated. If the previous link up notification was non-deterministic, then this notification includes an attribute signifying that the packets sent within the prior interval were lost.

If the bridge is identified as performing Rapid Spanning Tree Protocol (RSTP), it instead waits Bridge_Max_Age after packet reception (advertised in the BPDU's Max Age field), before forwarding. For ports which are known to be point-to-point through auto-negotiation, this delay is abbreviated to 3 seconds after auto-negotiation completes [IEEE-802.1D].

3.4.3. 802.1AB Link-Layer Discovery Protocol

The recently defined 802.1AB Link-Layer Discovery Protocol (LLDP) provides information to devices that are directly adjacent to them on the local LAN [IEEE-802.1ab].

LLDP sends information periodically and at link status change time to indicate the configuration parameters of the device. Devices may send or receive these messages, or do both.

The LLDP message may contain a System Capabilities TLV, which describes the MAC- and IP-layer functions that a device is currently using. Where a host receives the System Capabilities TLV indicating that no Bridging is occurring on the LLDP transmitter, no delays for STP calculation will be applied to packets sent through this transmitter. This would allow the generation of a link up notification.

Additionally, if a host receives a System Capabilities TLV indicating that the LLDP transmitter is a bridge, the host's advertisement that it is an (end-host) Station-Only may tell the bridge not to run STP and may immediately allow forwarding.

Proprietary extensions may also indicate that data forwarding is already available on such a port. Discussion of such optimizations is out of scope for this document.

Because the protocol is new and not widely deployed, it is unclear how this protocol will eventually affect DNA in IPv4 or IPv6 networks.

3.4.4. Other Heuristics

In 802.3 networks, Network Interface Cards (NICs) are often capable of returning a speed and duplex indication to the host. Changes in these characteristics may indicate a connection to a new layer 2 network.

3.4.5. Summary

Link-layer indications in Ethernet-like networks are complicated by additional unadvertised delays due to spanning tree calculations. This may cause re-indication or retraction of indications previously sent to upper layer protocols.

4. Security Considerations

Attackers may spoof various indications at the link layer, or manipulate the physical medium directly in an effort to confuse the host about the state of the link layer. For instance, attackers may spoof error messages or disturb the wireless medium to cause the host to move its connection elsewhere or even to disconnect. Attackers may also spoof information to make the host believe it has a connection when, in reality, it does not. In addition, wireless networks such as 802.11 are susceptible to an attack called the "Evil Twin" attack where an attacker sets up an Access Point with the same SSID as a legitimate one and gets the use to connect to the fake access point instead of the real one. These attacks may cause use of non-preferred networks or even denial of service.

This specification does not provide any protection of its own for the indications from the lower layers. But the vulnerabilities can be mitigated through the use of techniques in other parts of the protocol stack. In particular, it is recommended that authentication, replay, and integrity protection of link-layer management messages are enabled when available. For example, the IEEE 802.1ae standard [IEEE-802.1ae] defines such mechanisms for IEEE 802-compliant MAC layers. Additionally, the protocol stack may also use some network-layer mechanisms to achieve partial protection. For instance, SEND [RFC3971] could be used to confirm secure reachability with a router. However, network layer mechanisms are unable to deal with all problems, such as insecure lower-layer notifications that lead to the link not functioning properly.

5. Contributors

In addition to the people listed in the author list, text for the specific link-layer technologies covered by this document was contributed by Thomas Noel (IEEE 802.11b) and Greg Daley (IEEE 802.3). The authors would like to thank them for their efforts in bringing this document to fruition.

6. Acknowledgements

The authors would like to acknowledge Bernard Aboba, Sanjeev Athalye, JinHyeock Choi, John Loughney, Pekka Nikander, Brett Pentland, Tom Petch, Dan Romascanu, Pekka Savola, Steve Bellovin, Thomas Narten, Matt Mathis, Alfred Hoenes, and Muhammad Mukarram bin Tariq for their useful comments and suggestions.

7. References

7.1. Normative References

[CDMA2K] "cdma2000 Wireless IP Network Standard", , December 2000.

[GPRS] "Digital cellular telecommunications system (Phase 2+); General Packet Radio Service (GPRS) Service description; Stage 2", 3GPP TS 03.60 version 7.9.0 Release 98.

[GPRS-LINK] "Digital cellular telecommunications system (Phase 2+); Radio subsystem link control", 3GPP GSM 03.05 version 7.0.0 Release 98.

[IEEE-802.11a] Institute of Electrical and Electronics Engineers, "IEEE Std 802.11a-1999, supplement to IEEE Std 802.11-1999, Part 11: Wireless MAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: High-speed Physical Layer in the 5 GHZ band", IEEE Standard 802.11a, September 1999.

[IEEE-802.11b] Institute of Electrical and Electronics Engineers, "IEEE Std 802 Part 11, Information technology - Telecomunications 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 Standard 802.11b, August 1999.

[IEEE-802.11g] Institute of Electrical and Electronics Engineers, "IEEE Std 802.11g-2003, Amendment to IEEE Std 802.11, 1999 edition, Part 11: Wireless MAN Medium Access Control (MAC) and Physical Layer (PHY) specifications. Amendment 4: Further Higher Data Rate Extension in the 2.4 GHz Band", IEEE Standard 802.11g, June 2003.

[IEEE-802.11i] Institute of Electrical and Electronics Engineers, "Supplement to STANDARD FOR Telecommunications and Information Exchange between Systems - LAN/MAN Specific Requirements - Part 11: Wireless Medium Access Control (MAC) and physical layer (PHY) specifications: Specification for Enhanced Security", IEEE 802.11i, December 2004.

[IEEE-802.1D] Institute of Electrical and Electronics Engineers, "IEEE standard for local and metropolitan area networks - common specifications - Media access control (MAC) Bridges", ISO/IEC IEEE Std 802.1D, 2004.

[IEEE-802.1ab] Institute of Electrical and Electronics Engineers, "Draft Standard for Local and Metropolitan Networks: Station and Media Access Control Connectivity Discovery (Draft 13)", IEEE draft Std 802.1AB, 2004.

[IEEE-802.1ae] Institute of Electrical and Electronics Engineers, "IEEE Std 802.1AE, Local and Metropolitan Area Networks - Media Access Control (MAC) Security", IEEE Standard 802.1ae, June 2006.

[IEEE-802.3] Institute of Electrical and Electronics Engineers, "IEEE standard for local and metropolitan area networks - Specific Requirements, Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications", ISO/IEC IEEE Std 802.3, 2002.

[RFC1332] McGregor, G., "The PPP Internet Protocol Control Protocol (IPCP)", RFC 1332, May 1992.

[RFC1661] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.

[RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

[RFC2472] Haskin, D. and E. Allen, "IP Version 6 over PPP", RFC 2472, December 1998.

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

[RFC4135] Choi, JH. and G. Daley, "Goals of Detecting Network Attachment in IPv6", RFC 4135, August 2005.

7.2. Informative References

[GPRS-CN] "Technical Specification Group Core Network; Internetworking between the Public Land Mobile Network (PLMN) supporting packet based services and Packet Data Networks (PDN) (Release 6)", 3GPP TS 29.061 version 6.1.0 2004-06.

[GPRS-GSSA] "Technical Specification Group Services and System Aspect; General Packet Radio Service (GPRS) Service description; Stage 2 (Release 6)", 3GPP TS 23.060 version 6.5.0 2004-06.

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

[RFC4068] Koodli, R., "Fast Handovers for Mobile IPv6", RFC 4068, July 2005.

[RFC4881] El Malki, K., "Low-Latency Handoffs in Mobile IPv4", RFC 4881, June 2007.

Authors' Addresses

Suresh Krishnan (editor)
Ericsson Research
8400 Decarie Blvd.
Town of Mount Royal, QC
Canada

EMail: suresh.krishnan@ericsson.com

Nicolas Montavont
GET ENST Bretagne
2, rue de la chataigneraie
Cesson-Sevigne 35576
France

Phone: (33) 2 99 12 70 23
EMail: nicolas.montavont@enst-bretagne.fr

Eric Njedjou
France Telecom
4, Rue du Clos Courtel BP 91226
Cesson Sevigne 35512
France

Phone: +33 299124878
EMail: eric.njedjou@orange-ftgroup.com

Siva Veerepalli
Qualcomm
5775 Morehouse Drive
San Diego, CA 92131
USA

Phone: +1 858 658 4628
EMail: sivav@qualcomm.com

Alper E. Yegin (editor)
Samsung
Istanbul
Turkey

Phone: +90 533 348 2402
EMail: a.yegin@partner.samsung.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.