

Network Working Group  
INTERNET-DRAFT  
Category: Informational  
<[draft-iab-link-indications-00.txt](#)>  
[14](#) October 2004

B. Aboba, Ed.  
Internet Architecture Board  
IAB

## Architectural Implications of Link Layer Indications

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, and any of which I become aware will be disclosed, in accordance with [RFC 3668](#).

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>.

This Internet-Draft will expire on April 22, 2005.

### Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

### Abstract

As a performance optimization, proposals have been made for utilizing link layer indications (also known as "triggers" or "hints") to influence the behavior of the Internet, Transport or Application layers. This document briefly summarizes current proposals and describes architectural issues relating to link layer indications.

INTERNET-DRAFT

Link Layer Indications

14 October 2004

## Table of Contents

<a href="#">1.</a>	Introduction.....	<a href="#">3</a>
<a href="#">1.1</a>	Requirements .....	<a href="#">3</a>
<a href="#">1.2</a>	Terminology .....	<a href="#">3</a>
<a href="#">1.3</a>	Link Indications .....	<a href="#">5</a>
<a href="#">1.4</a>	Proposals .....	<a href="#">6</a>
<a href="#">1.5</a>	Layering Model .....	<a href="#">7</a>
<a href="#">1.6</a>	Link Behavior .....	<a href="#">10</a>
<a href="#">1.7</a>	Implementation Differences .....	<a href="#">11</a>
<a href="#">2.</a>	Architectural considerations .....	<a href="#">12</a>
<a href="#">2.1</a>	Model Validation .....	<a href="#">13</a>
<a href="#">2.2</a>	Robustness .....	<a href="#">15</a>
<a href="#">2.3</a>	Effectiveness .....	<a href="#">18</a>
<a href="#">2.4</a>	Interoperability Issues .....	<a href="#">19</a>
<a href="#">2.5</a>	Race Conditions .....	<a href="#">19</a>
<a href="#">2.6</a>	Layer Compression .....	<a href="#">23</a>
<a href="#">2.7</a>	Remoting Implications .....	<a href="#">24</a>
<a href="#">2.8</a>	Security Considerations .....	<a href="#">25</a>
<a href="#">3.</a>	Future Work .....	<a href="#">26</a>
<a href="#">4.</a>	References .....	<a href="#">27</a>
<a href="#">4.1</a>	Normative References .....	<a href="#">27</a>
<a href="#">4.2</a>	Informative References .....	<a href="#">27</a>
<a href="#">Appendix A</a>	- IAB Members .....	<a href="#">31</a>
	Intellectual Property Statement .....	<a href="#">31</a>
	Disclaimer of Validity .....	<a href="#">31</a>
	Copyright Statement .....	<a href="#">32</a>

INTERNET-DRAFT

Link Layer Indications

14 October 2004

## [1.](#) Introduction

As evidence mounts that correct utilization of link layer indications can provide real benefits, and that incorrect utilization can degrade performance, the importance of understanding the role of link indications in the Internet architecture has grown in importance.

This document is an attempt to summarize current understanding of the role of link layer indications, as well as to provide advice to document authors considering the role of link layer indications within their own work.

In [Section 1](#) of this document we present a brief overview of current proposals, as well as recent research on link behavior. Based on the overview, [Section 2](#) provides advice to document authors. [Section 3](#) describes future work.

### [1.1.](#) Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

### [1.2.](#) Terminology

#### Access Point (AP)

A station that provides access to the distribution services, via the wireless medium (WM) for associated stations.

#### Association

The service used to establish an access point/station (AP/STA) mapping and enable STA access to the Distribution System.

#### Basic Service Set (BSS)

A set of stations controlled by a single coordination function, where the coordination function may be centralized (e.g., in a

single AP) or distributed (e.g., for an ad-hoc network). The BSS can be thought of as the coverage area of a single AP.

#### Care of Address (CoA)

A unicast routable address associated with a mobile node while visiting a foreign link; the subnet prefix of this IP address is a foreign subnet prefix. Among the multiple care-of addresses that a mobile node may have at any given time (e.g., with different subnet prefixes), the one registered with the mobile node's home agent for a given home address is called its "primary" care-of address.

IAB

Informational

[Page 3]

---

INTERNET-DRAFT

Link Layer Indications

14 October 2004

#### Correspondent Node

A peer node with which a mobile node is communicating. The correspondent node may be either mobile or stationary.

#### Distribution System (DS)

A system used to interconnect a set of basic service sets (BSSs) and integrated local area networks (LANs) to create an extended service set (ESS).

#### Dynamic Host Configuration Protocol (DHCP) client

A DHCP client or "client" is an Internet host using DHCP [[RFC2131](#)] to obtain configuration parameters such as a network address.

#### DHCP server

A DHCP server or "server" is an Internet host that returns configuration parameters to DHCP clients.

#### Extended Service Set (ESS)

A set of one or more interconnected basic service sets (BSSs) and integrated local area networks (LANs) that appears as a single BSS to the logical link control layer at any station associated with one of those BSSs. The ESS can be thought of as the coverage area provided by a collection of APs all interconnected by the Distribution System. It may consist of one or more IP subnets.

#### Home Address (HoA)

A unicast routable address assigned to a mobile node, used as the permanent address of the mobile node. This address is within the mobile node's home link. Standard IP routing mechanisms will

deliver packets destined for a mobile node's home address to its home link. Mobile nodes can have multiple home addresses, for instance when there are multiple home prefixes on the home link.

#### Inter-Access Point Protocol (IAPP)

A protocol used between access points that assures that the station may only be connected to a single AP within the ESS at a time, and also provides for transfer of context to the new AP.

**Link** A communication facility or medium over which nodes can communicate at the link layer, such as an Ethernet (simple or bridged). The link layer is the layer immediately below IP.

#### Link Indication

Information provided by the link layer to higher layers relating to the state of the link.

#### Mobile Node

A node that can change its point of attachment from one link to

IAB

Informational

[Page 4]

---

INTERNET-DRAFT

Link Layer Indications

14 October 2004

another, while still being reachable via its home address.

#### Point of Attachment

A location within the network where a host may be connected. This attachment point can be characterized by its address prefix and next hop routing information.

#### Most Likely Point of Attachment (MLPA)

The point of attachment heuristically determined by the host to be most likely, based on hints from the network.

#### Routable address

In this specification, the term "routable address" refers to any address other than an IPv4 Link-Local address. This includes private addresses as specified in [[RFC1918](#)].

#### Station (STA)

Any device that contains an IEEE 802.11 conformant medium access control (MAC) and physical layer (PHY) interface to the wireless medium (WM).

#### Valid address

The term "valid address" refers to either a static IPv4 address, or an address assigned via DHCPv4 which has not been relinquished, and whose lease has not yet expired.

#### Weak End-System Model

In the Weak End-System Model, packets sent out an interface need not necessarily have a source address configured on that interface.

### [1.3.](#) Link Indications

A link indication represents information provided by the link layer to higher layers relating to the state of the link.

While link layer indications vary considerably between media, abstraction models have been proposed. For example, [\[GenTrig\]](#) defines "generic triggers", including "Link Up", "Link Down", "Link Going Down", "Link Going Up", "Link Quality Crosses Threshold", "Trigger Rollback", and "Better Signal Quality AP Available". Other link indications include the current link rate (which may vary with time and location), link identifiers (e.g. SSID, BSSID in 802.11), and statistics relating to link performance (such as the delay or loss rate).

Among the most commonly implemented link indications are the "Link Up" and "Link Down" indications, which are based on an idealized link behavior model originally developed for wired networks. This model

assumes that links in the "Up" state experience low frame loss in both directions and are ready to send and receive IP data packets. Similarly, it is assumed that a link that is in the "Down" state is unsuitable for sending and receiving IP data packets in either direction.

Link indications based on signal quality, such as "Link Going Down", "Link Going Up", and "Link Quality Crosses Threshold" are primarily intended for use in handoff optimization. These indications assume an idealized model of radio propagation, where signal strength varies smoothly and frame loss is well predicted by signal strength and distance.

### [1.4.](#) Proposals

Within the Internet Layer, proposals have been made for utilizing link layer indications to optimize IP configuration, to improve the usefulness of routing metrics, and to optimize aspects of Mobile IP handoff.

In "Detection of Network Attachment (DNA) in IPv4" [[DNAv4](#)], link layer indications are utilized to optimize Internet layer configuration. This enables a host that has moved to a new point of attachment but remained within the same subnet to rapidly confirm a currently valid configuration, rather than utilizing the DHCP protocol [[RFC2131](#)].

"A High-Throughput Path Metric for Multi-Hop Wireless Routing" [[ETX](#)] describes how routing metrics can be improved by utilizing the Expected Transmission Count (ETX) metric, which takes link layer loss rates into account, enabling the selection of routes maximizing available throughput. While the ETX metric did not take the negotiated rate into account, this was noted as a subject for further study.

In "L2 Triggers Optimized Mobile IPv6 Vertical Handover: The 802.11/GPRS Example" [[Park](#)] the authors propose that the mobile node send a router solicitation on receipt of a "Link Up" indication in order provide lower handoff latency than would be possible using generic movement detection [[RFC3775](#)]. The authors also suggest immediate invalidation of the Care-Of-Address (CoA) on receipt of a "Link Down" indication.

Within the Transport Layer, proposals have focused on countering the effects of handoff-induced packet loss. This includes proposals for improving transport parameter estimation, as well as triggering immediate retransmission on availability of an interface or intervening link.

"Framework and Requirements for TRIGTRAN" [[TRIGTRAN](#)] discusses optimizations to recover earlier from a retransmission timeout incurred during a period in which an interface or intervening link was down.

"Link-layer Triggers Protocol" [[Yegin](#)] describes transport issues arising from lack of host awareness of link conditions on downstream Access Points and routers. A link-layer trigger remoting is proposed

to address the issue.

In "TCP Extensions for Immediate Retransmissions" [[Eggert](#)], it is proposed that in addition to regularly scheduled retransmissions that retransmission be attempted by the transport layer on receipt of an indication that connectivity to a peer node may have been restored. End-to-end connectivity restoration indications include "Link Up", confirmation of first-hop router reachability, confirmation of Internet layer configuration, and receipt of other traffic from the peer.

In "The BU-trigger method for improving TCP performance over Mobile IPv6" [[Kim](#)], the authors note that handoff-related packet loss is interpreted as congestion by the transport layer. In the case where the correspondent node is sending to the mobile node, it is proposed that receipt of a Binding Update by the correspondent node be used as a signal to the transport layer to adjust cwnd and ssthresh values, which may have been reduced due to handoff-induced packet loss. The authors recommend that cwnd and ssthresh be recovered to pre-timeout values, regardless of whether the link parameters have changed. The paper does not discuss the behavior of a mobile node sending a Binding Update, in the case where the mobile node is sending to the correspondent node.

At the application layer, the usage of "Link Down" indications has been proposed to augment presence systems. In such systems, client devices periodically refresh their presence state using application layer protocols such as SIMPLE [[RFC3428](#)] or XMPP [[RFC3921](#)]. If the client should become disconnected, their unavailability will not be detected until the presence status times out, which can take many minutes. However, if a link goes down, and a disconnect indication can be sent to the presence server (presumably by the access point, which remains connected), the status of the user's communication application can be updated nearly instantaneously.

### [1.5.](#) Layering Model

A simplified layered indication model is shown in Figure 1. This model includes both internally generated link indications as well as Internet and Transport layer indications arising out of external

interactions (such as receipt of Mobile IP Binding Updates, and



detection of path changes via routing protocols and TTL changes).

In this model, link indications provided to higher layers include the frame loss rate (before retransmissions), the current link rate, the link state (up/down), and link identifiers. These link indications are inter-dependent. For example, the rate adjustment and detection algorithms are typically influenced by frame loss, and in turn, the determination of a "Link Down" indication may be influenced by the detection and search process. Link Identifiers are typically obtained in the process of bringing the link to the "Up" state.

Link indications may be utilized by the Internet layer in order to optimize aspects of IP configuration, routing and mobility. As noted in [\[DNAv4\]](#), "Link Up" indications and link Identifiers may be useful in validation of an existing IP configuration. Once the IP configuration is confirmed, it may be determined that an IP address change has occurred.

As described in [\[ETX\]](#), the frame loss rate as well as the current link rate may be utilized in the calculation of routing metrics. Within "Weak End-System Model" implementations, changes in routing metrics may in turn result in a change in the outgoing interface for one or more transport connections. Routes may also be added or withdrawn, resulting in loss or gain of peer connectivity. The Internet layer may also become aware of path changes, by other mechanisms such as a change in the IP TTL of received packets.

A change in the outgoing interface may in turn influence the mobility sub-layer if present, causing a change in the incoming interface. The mobility sub-layer may also become aware of a change in the incoming interface of a peer (via receipt of a Mobile IP binding update).

Internet layer indications such as IP address and path changes are provided to the Transport layer, which also receives Link layer indications such as the loss rate, and "Link Up"/"Link Down".

Based on these Internet layer indications, the transport layer may wish to modify transport parameter estimates (such as by resetting parameter estimates of connections undergoing a path change), or tear down transport connections (due to invalidation of a connection's originating IP address). It is also possible for the transport layer to utilize "Link Up"/"Link Down" indications, and Loss Rate information to improve transport parameter estimates. However, the required algorithms not well understood.

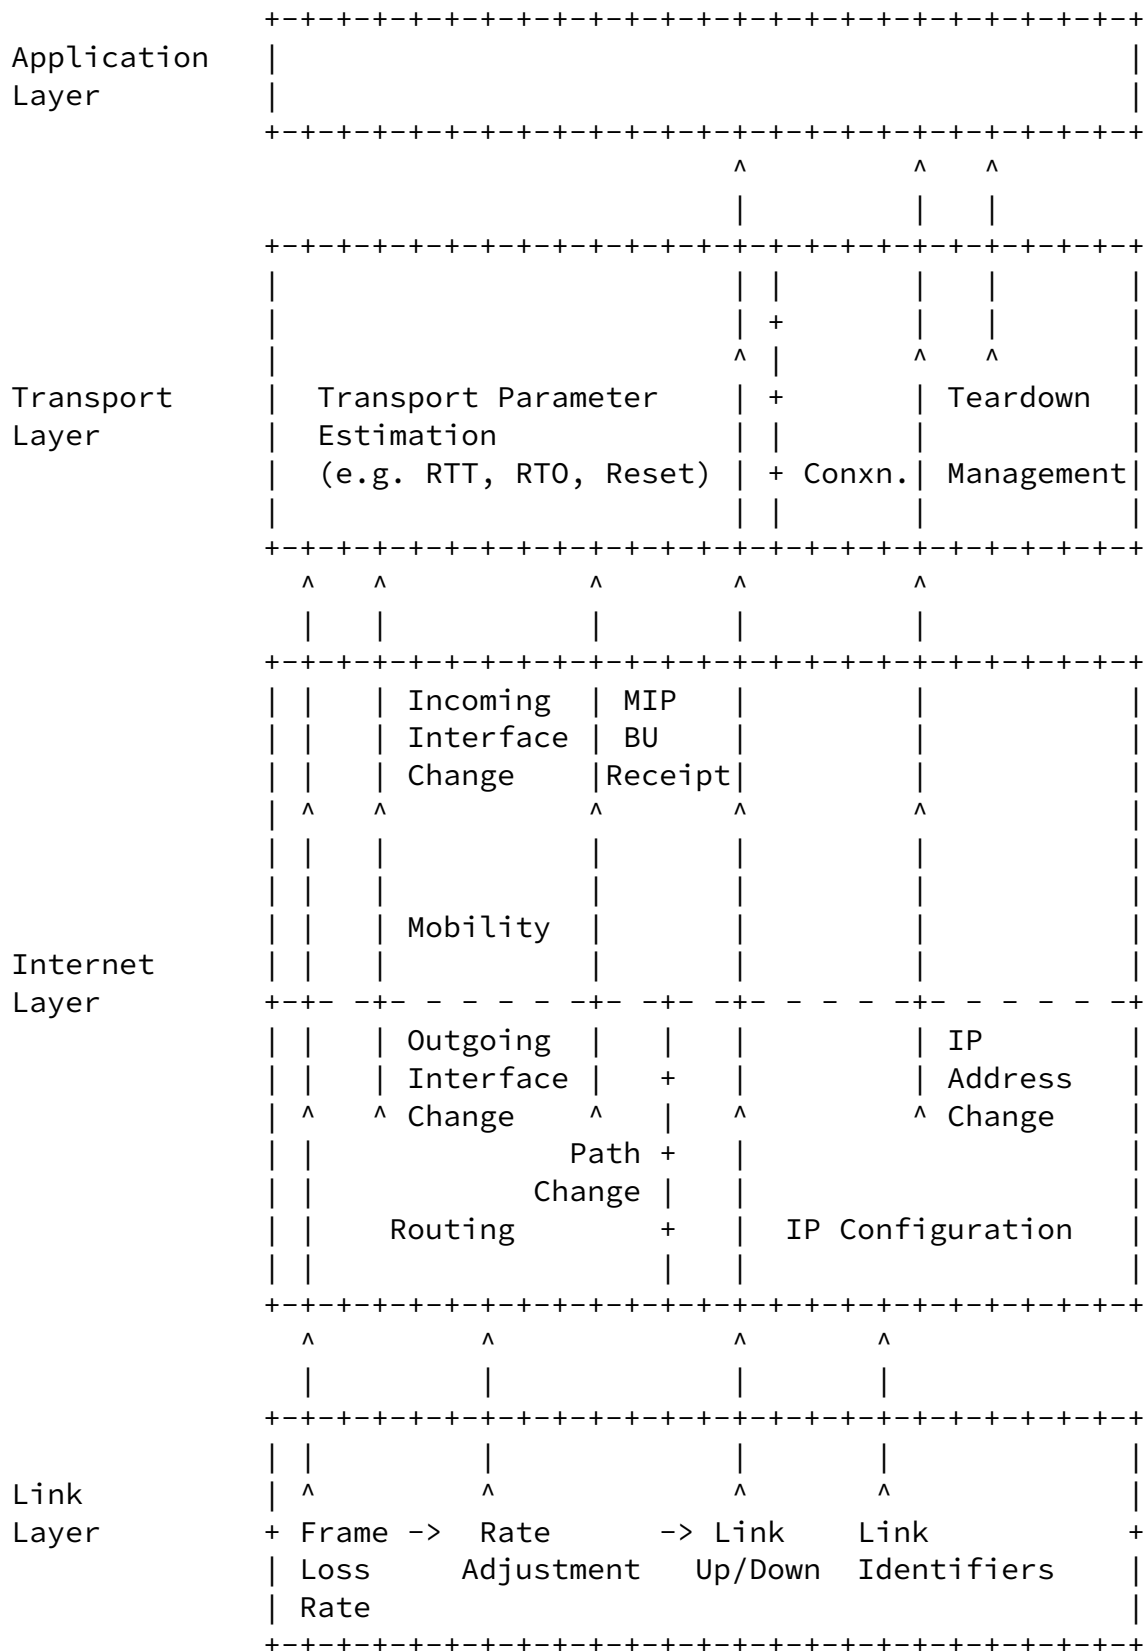


Figure 1. Layered Indication Model

INTERNET-DRAFT

Link Layer Indications

14 October 2004

In addition to Internet layer indications propagated to the Application layer (such as IP address changes), the Transport layer propagates its own indications, such as connection teardown. In most cases applications can obtain the information they need from Internet and Transport layer indications so that they do not need to directly consume link indications.

#### [1.6.](#) Link Behavior

In order to understand the applicability of the layered indication model it is instructive to review recent research relating to link performance.

In "Measurement and Analysis of the Error Characteristics of an In-Building Wireless Network" [[Eckhardt](#)], the authors characterize the performance of an AT&T Wavelan 2 Mbps in-building WLAN operating in Infrastructure mode on the Carnegie-Mellon Campus. In this study, very low frame loss was experienced. As a result, links could either be assumed to operate very well or not at all.

In "Performance of Multihop Wireless Networks: Shortest Path is Not Enough" [[Shortest](#)] the authors studied the performance of both an indoor and outdoor mesh network. By measuring inter-node throughput, the best path between nodes was computed. The throughput of the best path was compared with the throughput of the shortest path computed based on a hop-count metric. In almost all cases, the shortest path route offered considerably lower throughput than the best path.

In examining link behavior, the authors found that rather than exhibiting a bi-modal distribution between "up" (low loss rate) and "down" (high loss rates), many links exhibited intermediate loss rates. Asymmetry was also common, with 30 percent of links demonstrating substantial differences between in the loss rates in each direction. As a result, on wireless networks the measured throughput can differ substantially from the negotiated rate due to retransmissions, and successful delivery of routing packets is not necessarily an indication that the link is useful for delivery of data.

"Link-level Measurements from an 802.11b Mesh Network" [[Aguayo](#)] analyzes the causes of frame loss in a 38-node urban multi-hop 802.11 ad-hoc network. In most cases, links that are very bad in one direction tend to be bad in both directions, and links that are very good in one direction tend to be good in both directions. However, 30 percent of links exhibited loss rates differing substantially in each direction.

Signal to noise ratio and distance showed little value in predicting

loss rates, and rather than exhibiting a step-function transition between "up" (low loss) or "down" (high loss) states, inter-node loss rates varied widely, demonstrating a nearly uniform distribution over the range at the lower rates. The authors attribute the observed effects to multi-path fading, rather than attenuation or interference.

The findings of [[Eckhardt](#)] and [[Aguayo](#)] demonstrate the diversity of loss conditions observed in practice. There is a fundamental difference between infrastructure networks in which site surveys and careful measurement can assist in promoting ideal behavior and ad-hoc/mesh networks in which node mobility and external factors such as weather may not be easily controlled.

### [1.7](#). Implementation Differences

The literature also describes the effect of implementation differences on link indications. For the purposes of illustration, we will restrict ourselves to literature relating to IEEE 802.11 implementations.

"An Empirical Analysis of the IEEE 802.11 MAC Layer Handoff Process" [[Mishra](#)] investigates handoff latencies obtained with three mobile STAs implementations communicating with two APs. The study found that there is large variation in handoff latency among STA and AP implementations and that implementations utilize different message sequences. For example, one STA sends a Reassociation Request prior to authentication, which results in receipt of a Deauthenticate message. The study divided handoff latency into discovery, authentication and reassociation exchanges, concluding that the discovery phase was the dominant component of handoff delay. Detection was not investigated.

"Techniques to reduce IEEE 802.11b MAC layer handover time" [[Velayos](#)] measured handover times for a stationary STA after the AP was turned off. This study divided handover times into detection (determination of the need for handover), search (discovery of alternative attachment points), and execution phases (authentication and association exchanges). These measurements indicated that the detection phase was longest in duration. The duration of the detection phase is determined by the number of non-acknowledged frames triggering the search phase and precursors such as RTS/CTS and rate adaptation.

Detection behavior varied widely between implementations. For example, NICs designed for desktops attempted more retransmissions prior to triggering search as compared with laptop designs, since they assumed that the AP is always in range, regardless of Beacon

reception.

The study recommends that the duration of the detection phase be reduced by initiating the search phase as soon as collisions can be excluded as the cause of non-acknowledged transmissions; the authors recommend three consecutive transmission failures as the cutoff. Where the STA is not sending, it is recommended that Beacon reception be tracked if no data frames are being received, and that Beacon spacing be reduced to 60 ms in order to reduce detection times. In order to compensate for more frequent triggering of the search phase, the authors recommend algorithms for wait time reduction, as well as interleaving of search with data transmission.

"Roaming Interval Measurements" [[Alimian](#)] presents data on stationary STAs after the AP signal has been shut off. This study highlighted implementation differences in rate adaptation as well as detection, scanning and handoff. As in [[Velayos](#)], performance varied widely between implementations, from half an order variation in rate adaptation to an order of magnitude difference in connectivity detection times, two orders of magnitude in scanning, and one and a half orders of magnitude in handoff times.

"An experimental study of IEEE 802.11b handoff performance and its effect on voice traffic" [[Vatn](#)] describes handover behavior observed when the signal from AP is gradually attenuated, which is more

representative of field experience than the shutoff techniques used in [Velayos]. Stations were configured to initiate handover when signal strength dipped below a threshold, rather than purely based on frame loss, so that they could begin handover while still connected to the current AP. It was noted that stations continue to receive data frames during the search phase. Station-initiated Disassociation and pre-authentication were not observed in this study.

## [2.](#) Architectural considerations

While the literature on the usage of link layer indications provides persuasive evidence of their utility, experience shows that a number of difficulties can arise in making effective use of them. These issues include:

- a. Model validation
- b. Robustness
- c. Effectiveness
- d. Interoperability Issues
- e. Race conditions
- f. Layer compression
- g. Remoting implications

### h. Security implications

The sections that follow discuss each of these issues in turn.

#### [2.1.](#) Model Validation

In "The mistaken axioms of wireless-network research" [Kotz], the authors conclude that mistaken assumptions relating to link performance may lead to the design of network protocols that may not work in practice. In order to avoid these pitfalls, documents dependent on link indications should explicitly articulate the assumptions of the link model and describe the circumstances in which it applies.

Authors need to be careful to avoid use of simplified network models in circumstances where the model does not apply. For example, generic "trigger" models assume that a link is either in a state experiencing low frame loss ("Link Up") or in a state where few if

any frames are delivered ("Link Down"). Often symmetry is assumed as well, so that a link is assumed to be either "Up" in both directions or "Down" in both directions. In wireless networks, particularly in the case of ad-hoc or mesh deployments, these assumptions may prove invalid.

Furthermore, where links are in intermediate states between "Up" and "Down" and asymmetry is encountered, generic "triggers" such as "Link Going Down", "Link Going Up", "Link Quality Crosses Threshold" may prove difficult to define and may prove to be unreliable predictors of future link performance.

Once the network model is defined, considerable effort may be required to map the model to common link types. In practice the definition of "Link Up" or "Link Down" may vary according to the link layer. For example, within PPP [[RFC1661](#)], either peer may send an LCP-Terminate frame in order to terminate the PPP link layer, and a link may only be assumed to be usable for sending network protocol packets once NCP negotiation has completed for that protocol.

Within IEEE 802.11, the definition of "Link Up" and "Link Down" depends on whether the station is mobile or stationary, whether infrastructure or ad-hoc mode is in use, and whether security and Inter-Access Point Protocol (IAPP) is implemented.

Where a mobile 802.11 STA encounters a series of consecutive non-acknowledged frames, the most likely cause is that the station has moved out of range of the AP. As a result, [[Velayos](#)] recommends that the station begin the search phase after collisions can be ruled out, after three consecutive non-acknowledged frames. Only when no

alternative point of attachment is found is a "Link Down" indication returned.

In a stationary point-to-point installation, the most likely cause of an outage is that the link has become impaired. As a result, implementations tend to be more persistent and a "Link Down" indication may be returned later.

In Infrastructure mode, IEEE 802.11-2003 enables reception of data frames only in State 3 ("Authenticated" and "Associated"). As a result, a transition to State 3 (e.g. completion of a successful

Association or Reassociation exchange) enables sending and receiving of network protocol packets and a transition from State 3 to State 2 (reception of a "Disassociate" frame) or State 1 (reception of a "De-authenticate" frame) disables sending and receiving of network protocol packets. As a result, IEEE 802.11 stations typically signal "Link Up" on receipt of a successful Association or Reassociation response.

Within the [[IEEE80211f](#)] specification, after sending a Reassociation Response, an Access Point will send a frame with the station's source address to a multicast destination. This causes switches within the Distribution System (DS) to update their learning tables, readying the DS to forward frames to the station at its new point of attachment. Were the AP to not send this "spoofed" frame, the station's location would not be updated within the DS until it sent its first frame at the new location. Thus IAPP serves to equalize uplink and downlink handover times.

The signalling of "Link Down" is considerably more complex. Even though a transition to State 2 or State 1 results in the station being unable to send or receive IP packets, this does not necessarily imply that such a transition should be considered a "Link Down" indication. In an infrastructure network, a station may have a choice of multiple access points offering connection to the same network. In such an environment, a station that is unable to reach State 3 with one access point may instead choose to attach to another access point. Rather than registering a "Link Down" indication with each move, the station may instead register a series of "Link Up" indications.

In [[IEEE80211i](#)] forwarding of frames from the station to the distribution system is only feasible after the completion of the 4-way handshake and group-key handshake, so that entering State 3 is no longer sufficient.

Unfortunately, other elements of the IEEE 802.11 specification such as IEEE 802.11f continue to recognize the Reassociation Response as

the "Link Up" definition. By spoofing a multicast frame with the station's source address once it sends a Reassociation Response, Access Points implementing IEEE 802.11f cause the learning tables within switches comprising the DS to be updated. This enables an



attacker to deny service to attached stations by sending a Reassociation Request from anywhere within the ESS. Without the spoofing recommended in IEEE 802.11f, such an attack would only be able to disassociate stations on the AP to which the Reassociation Request was sent.

[IEEE80211i] implementations utilizing the "Link Up" definition from [IEEE80211] or [IEEE80211f] have also encountered difficulty in IP address assignment, since they may trigger DHCP [RFC2131] or RS/RA prior to when the link is usable by the Internet layer. As a result, Internet layer configuration may fail.

In contrast, in 802.11 ad-hoc mode with no security, reception of data frames is enabled in State 1 ("Unauthenticated" and "Un-associated"). As a result, reception of data frames is enabled at any time, and no explicit "Link Up" indication exists.

## [2.2.](#) Robustness

Implementation experience provides us with several examples of situations in which improper consideration of link layer indications can result in operational malfunctions. Given the potential problems, proposals for consideration of link layer indications must demonstrate robustness against misleading indications. Elements of robustness include:

- a. Indication validation
- b. Damping and hysteresis

### [2.2.1.](#) Indication Validation

As noted in [Section 1.6](#) and 1.7, radio propagation and implementation differences can impact the reliability of link layer indications. [Kotz] notes that the three-dimensional nature of wireless propagation can result in large signal strength changes over short distances, generating short-lived "Link Down" and "Link Up" indications that are not be predicted by a two dimensional radio propagation model.

As described in [Aguayo], wireless links often exhibit loss rates intermediate between "up" (low loss) and "down" (high loss) states, as well as substantial asymmetry. In these circumstances, a "Link Up" indication may not imply bi-directional reachability. Also, a reachability demonstration based on small packets may not mean that

the link is suitable for carrying larger data packets. As a result, "Link Up" and "Link Down" indications may not reliably determine whether a link is suitable for carrying IP traffic.

Where the reliability of a link layer indication is suspect, it is best to treat the indication as a "hint" that is advisory in nature, rather than a "trigger" forcing a given action. In order to provide increased robustness, heuristics can be developed to determine whether the "hint" is valid or should be discarded.

In addition, a recovery step may be utilized in order to limit the potential damage from link indications determined to be invalid after they have been acted on.

To provide robustness in the face of potentially misleading link indications, in [\[DNav4\]](#) "Link Up" indications are assumed to be inherently unreliable, so that bi-directional reachability needs to be demonstrated prior to validating an existing IP configuration. However, in the case of a link of intermediate loss rate, success with the [\[DNav4\]](#) reachability test does not guarantee that the link is suitable for carrying data.

Another example of link indication validation occurs in IPv4 Link-Local address configuration. Prior to configuration of an IPv4 Link-Local address, it is necessary to run a claim and defend protocol [\[RFC3927\]](#). Since a host needs to be present to defend its address against another claimant, and address conflicts are relatively likely, a host returning from sleep mode or receiving a "Link Up" indication could encounter an address conflict were it to utilize a formerly configured Link-Local Link-Local address without rerunning claim and defend.

### [2.2.2](#). Damping and Hysterisis

Damping and hysterisis can be utilized to ensure that stability is maintained in the face of jittery link indications. These limits typically place constraints on the number of times a given action can be performed within a time period or introduce damping mechanisms to prevent instability.

While [\[Aguayo\]](#) found that frame loss was relatively stable for stationary stations, obstacles to radio propagation and multipath interference can result in rapid changes in signal strength for a mobile station. As a result, it is possible for mobile stations to encounter rapid changes in link performance, including changes in the negotiated rate, frame loss and even "Link Up"/"Link Down" indications.

INTERNET-DRAFT

Link Layer Indications

14 October 2004

Where link-aware routing metrics are implemented, this can result in rapid metric changes, potentially resulting in frequent changes in the outgoing interface for "Weak End-System" implementations. As a result, it may be necessary to introduce route flap dampening.

However, the benefits of damping need to be weighed against the additional latency that can be introduced. For example, in order to filter out spurious "Link Down" indications, these indications may be delayed until it can be determined that a "Link Up" indication will not follow shortly thereafter. However, in situations where multiple Beacons are missed such a delay may not be needed, since there is no evidence of a suitable point of attachment in the vicinity.

In some cases, it may be desirable to ignore link indications entirely. Since it is possible for a host to transition from an ad-hoc network to a network with centralized address management, a host receiving a "Link Up" indication cannot necessarily conclude that it is appropriate to configure a IPv4 Link-Local address at all.

It can be argued that reliable transport protocol implementations should ignore "Link Down" indications, rather than tearing down connections, regardless of the cause of the "Link Down" indication.

Where the "Link Down" indication results from frame loss rather than an explicit exchange, the indication may be transient, rapidly followed by a "Link Up" indication. Even where the "Link Down" indication results from an explicit exchange such as a PPP LCP-Terminate or an 802.11 Disassociation or Deauthenticate, an alternative point of attachment may be available, allowing connectivity to be quickly restored. As a result, robustness is best served by allowing connections to remain up until the connection source address is invalidated by an address change, or the connection times out.

Where link indications are used to optimize transport performance, authors must demonstrate that effective congestion control is maintained [[RFC2914](#)] in the face of rapidly changing link indications.

In addition, where a proposal involves "recovery" from a handoff

event, it is important to demonstrate that the recovered parameters (such as the adjusted RTT, RTO, congestion window, etc.) remain valid, as noted in [[RFC2861](#)].

Consider a proposal where a "Link Up" indication is used by a router to signal retransmission a previously sent packet, in order to enable ACK reception prior to expiration of the host's retransmission timer. Where "Link Up" indications follow in rapid succession, this could

result in a burst of retransmitted packets, violating the law of conservation of packets.

At the Application Layer, link layer indications have been utilized by applications such as Presence [[RFC2778](#)] in order to optimize registration and user interface update operations. For example, implementations may attempt presence registration on receipt of a "Link Up" indication, and presence deregistration by a surrogate receiving a "Link Down" indication. Presence implementations using "Link Up"/"Link Down" indications this way violate the principle of "conservation of packets" when link indications are generated on a time scale of RTO or less. The problem is magnified since for each presence update, notifications can be delivered to many watchers.

The issue can be addressed by one or more of the following techniques:

- [a] Rate limiting. A limit of one packet per RTO can be imposed on packets generated from receipt of link indications.
- [b] Utilization of upper layer indications. Instead of consuming a "Link Up" indication, applications can consume alternative upper layer indications such as an IP address change notification.
- [c] Keepalives. Instead of consuming a "Link Down" indication, an application can utilize an application keepalive or consume transport layer indications such as connection teardown.

### [2.3.](#) Effectiveness

While link layer indications may show promise, it may be difficult to prove that processing of a given indication provides benefits in a wide variety of circumstances. Where link layer indications are

utilized for the purpose of optimization, proposals need to carefully analyze the effectiveness of the optimizations in the face of unreliable link layer indications. Since optimizations typically bring with them increased complexity, an optimization that does not bring about a performance improvement is not useful.

As with any optimization, the usefulness of link layer indications lies in demonstrated effectiveness of the optimization under consideration. This in turn may depend heavily on the penalty to be paid for false positives and false negatives.

As noted in [[DNAv4](#)], it is simultaneously possible for a link layer indication to be highly reliable, as well as for that indication to provide no net benefit, depending on the probability of a false indication and the penalty paid for the false indication.

In the case of [[DNAv4](#)], the benefits of successful optimization are modest, but the penalty for falsely concluding that the subnet remains unchanged is a lengthy timeout. The result is that link layer indications may not be worth considering if they are incorrect even just a small fraction of the time.

For example, it can be argued that a change in the Service Set Identifier (SSID) in [[IEEE80211](#)] is not a sufficiently reliable indication of subnet change. Within IEEE 802.11, the Service Set Identifier (SSID) functions as a non-unique identifier of the administrative domain of a Wireless LAN. Since the SSID is non-unique, many different operators may share the same SSID, and Access Points typically ship with a default value for the SSID (e.g. "default"). Since the SSID relates to the administrative domain and not the network topology, multiple SSIDs may provide access to the same prefix, and a single SSID may provide access to multiple prefixes at one or multiple locations.

Given this, it is unreliable to use the SSID alone for the purpose of movement detection. A host moving from one point of attachment to another, both with the same SSID, may have remained within the same subnet, or may have changed subnets. Similarly, a host discovering that the SSID has changed may have changed subnets, or it may not have. Moreover, where private address space is in use, it is possible for the SSID, the prefix (e.g. 192.168/16) and even the default gateway IP address to remain unchanged, yet for the host to

have moved to a different point of attachment. Were the host to make decisions relating to configuration of the IP layer (such as address assignment) based solely on the SSID, address conflicts are likely.

#### [2.4.](#) Interoperability Issues

Since link layer indications are often processed by upper layers for the purpose of optimization, proposals must demonstrate that interoperability remains possible (though potentially with degraded performance) even if one or more participants do not implement the proposals.

Where link layer indications are proposed for use in optimizing configuration of the Internet layer, it is necessary to demonstrate that the proposal does not interfere with routing protocol behavior, make address collisions more likely, or compromise Duplicate Address Detection (DAD).

#### [2.5.](#) Race Conditions

It is possible for link layer indications to be utilized directly by multiple layers of the stack in situations in which strict layering

may not be observed. In these situations, it is possible for race conditions to occur.

For example, as discussed earlier, link layer indications have been shown to be useful in optimizing aspects of Internet Protocol layer addressing and configuration as well as routing. Although [\[Kim\]](#) describes situations in which link layer indications are first processed by the Internet Protocol layer (e.g. MIPv6) before being consumed by the Transport Layer, in some situations it may be desirable for the Transport Layer to consume link layer indications directly.

For example, in situations where the "Weak End-System Model" is implemented, a change of outgoing interface may occur at the same time the Transport Layer is modifying transport parameters based on other link layer indications. As a result, transport behavior may differ depending on the order in which the link indications are processed.

When a multi-homed host experiences high frame loss on one of its interfaces, the ETX metric computed for that interface will rise, causing a change in the outgoing interface for one or more transport connections. This may trigger Mobile IP signaling so as to cause a change in the incoming path as well. At the same time, the Transport Layer may be estimating transport parameters based on the former outgoing interface, and may not properly adjust for the changes in outgoing and incoming paths.

To avoid race conditions, the following measures are recommended:

- a. Path change processing
- b. Layering
- c. Metric consistency

#### [2.5.1.](#) Path Change Processing

When the Internet layer detects a path change, such as a change in the outgoing or incoming interface of the host or the incoming interface of a peer, or perhaps a substantial change in the TTL of received IP packets, it may be worth considering whether to reset transport parameters to their initial values and allow them to be re-estimated. This ensures that estimates based on the former path do not persist after they have become invalid.

#### [2.5.2.](#) Layering

Another technique to avoid race conditions is to rely on layering to damp potentially misleading indications and provide greater link

layer independence.

The Internet layer is responsible for routing as well as IP configuration, and mobility, providing higher layers with an abstraction that is independent of link layer technologies. Since one of the major objectives of the Internet layer is maintaining link layer independence, upper layers relying on Internet Layer indications rather than consuming link layer indications directly can avoid link layer dependencies.

For example, in order to provide robustness, it is necessary to demonstrate that a link providing a "Link Up" indication is likely to

be usable for the transmission of IP data packets prior to using it. While applications can in principle incorporate their own versions of such a test, efficiency and maintenance considerations argue for providing this facility within the Internet layer.

Many "Link Up" indications do not result in a change of Internet Layer configuration, and many changes in link rate or frame loss do not result in a change of outgoing interface. By filtering "Link Up" indications, and selecting outgoing and incoming interfaces based on the link rate and frame loss, the Internet Layer enables upper layers to avoid writing their own code to filter and validate link indications.

The transport layer consumes Internet layer indication such as changes in the incoming/outgoing interface and Internet layer configuration changes, as well as potentially utilizing link layer indications directly. For example, the Internet layer may receive a "Link Down" indication followed by a subsequent "Link Up" indication. This information may be of interest to the Transport layer even if the Internet layer configuration does not change, since it may provide information about the validity of the transport parameters estimates.

In general, it is advisable for applications to consume indications from the Internet or Transport layers rather than consuming link indications directly, since this enables applications to leverage layered indication processing, resulting in improved robustness and scalability. For example, instead of directly consuming "Link Down" indications, applications may wish to rely on Application layer keepalives or the Transport layer to determine whether connections should be torn down; instead of consuming "Link Up" indications, applications can consume IP address change indications.

### [2.5.3.](#) Metric Consistency

Once a link is in the "Up" state, its effectiveness in transmission of data packets can be determined. For example, frame loss may be used in rate adjustment and detection of when to roam to an



alternative point of attachment. While connected, the effective throughput may be determined based on the negotiated rate and frame loss, and used in calculation of the routing metric, as described in [ETX].

However, prior to sending data packets over the link, other measures of suitability are required. As noted in [Shortest], the ability of a link to successfully transmit short frames utilized for control, management or routing is not indicative of its usefulness in carrying IP data packets. As a result, it is typically not possible to predict the negotiated rate or data frame loss rate in advance of a roaming decision.

As a result, 802.11 stations evaluating the suitability of candidate APs often utilize received signal strength and/or AP load as a primary selection criteria. Similarly, in order to enable stations to roam prior to encountering packet loss, studies such as [Vatn] have suggested using signal strength as a detection mechanism, rather than frame loss, as suggested in [Velayos].

The "Link Going Down", "Link Going Up", "Link Quality Crosses Threshold" indications were developed primarily to assist with handoff between interfaces, and are oriented toward inferred rather than measured suitability.

Research indicates that this approach may have some promise. For example, [Vertical] proposes use of signal strength and link utilization in order to optimize vertical handoff and demonstrates improved TCP throughput. However, without careful design, potential differences between link indications used in routing and those used in roaming and/or link enablement can result in instability, particularly in multi-homed hosts.

For example, receipt of "Link Going Down" or "Link Quality Crosses Threshold" indications could be used as a signal to enable another interface. However, unless the new interface is the preferred route for one or more destination prefixes, a "Weak End-System" implementation will not use the new interface for outgoing traffic. Where "idle timeout" functionality is implemented, the unused interface will be brought down, only to be brought up again by the link enablement algorithm.

As noted in [Aguayo], signal strength and distance are not good

predictors of frame loss or negotiated rate, due to the potential effects of multi-path interference. As a result a link brought up due to good signal strength may subsequently exhibit significant frame loss, and a low negotiated rate. Similarly, an AP demonstrating low utilization may not necessarily be the best choice, since utilization may be low due to hardware or software problems. As noted in [[Villamizar](#)], link utilization-based routing metrics have a history of instability, so that they are rarely deployed.

## [2.6.](#) Layer compression

In many situations, the exchanges required for a host to complete a handoff and reestablish connectivity are considerable. This includes link layer scanning, authentication and connectivity establishment; Internet layer configuration, routing and mobility exchanges; transport layer retransmission and recovery; security association re-establishment; application protocol reauthentication and reregistration exchanges, etc. Given this, it is natural to consider combining exchanges occurring within multiple layers into a single exchange.

Often this combined exchange occurs within the link layer. For example, in [[EAPoL](#)], a link layer EAP exchange may be used for the purpose of IP address assignment, potentially bypassing Internet layer configuration. Within [[PEAP](#)], it is proposed that a link layer EAP exchange be used for the purpose of carrying Mobile IPv6 Binding Updates. [[MIPEAP](#)] proposes that EAP exchanges be used for configuration of Mobile IPv6.

While the goals of layer compression are laudable, care needs to be taken to avoid compromising interoperability and introducing link layer dependencies into the Internet and Transport layers. For example, where Link layer and Internet or Transport layer mechanisms are combined, it is necessary for hosts to maintain the ability to interoperate without layer compression schemes, in order to permit operation on networks where they are not available.

As noted earlier, while the layered handling of link indications introduces latency, it also increases robustness. As a result, layer compression schemes need to take care to avoid introducing unnecessary brittleness. For example, in order to optimize IP address assignment, it has been proposed that prefixes be advertised at the link layer. While in theory such a proposal addresses the non-uniqueness issues found with the use of the SSID for movement detection, it suffers from its own set of problems.

For example, [[IEEE8021X](#)] enables the VLANID to be assigned dynamically. As a result, a prefix advertised at the link layer need

INTERNET-DRAFT

Link Layer Indications

14 October 2004

not correspond to the prefix assigned to the host once it connects to the link. Were IP configuration to be based on such hints, errors are likely.

### [2.7.](#) Remoting implications

Proposals which include support for remoting of link layer indications need to carefully consider the layering, security and transport implications.

While facilities such as ICMP "source quench" were originally provided at the Internet layer, these facilities have fallen into disuse due to their questionable value for the Transport layer. In general, the Transport layer is able to determine an appropriate (and conservative) response to congestion based on packet loss or explicit congestion notification, so that ICMP "source quench" indications are not needed, and in fact the sending of additional "source quench" packets during periods of congestion may be detrimental.

On the other hand, proposals such as [\[ETX\]](#) imply that hosts participating in the routing mesh may gain knowledge of remote link conditions where link-aware routing metrics are used. This can be accomplished securely if routing protocol security is implemented.

For example, when a link experiences frame loss, the ETX metric will increase, possibly resulting in selection of an alternate route. If the troubled link represents the only path to a prefix and the link experiences high frame loss ("Down"), the route will be withdrawn or the metric will become infinite. Thus, where link-indication aware routing metrics are implemented, "link indication remoting" can be accomplished via host participation in the routing mesh, and transport layer response to path changes.

Proposals involving remoting of link layer indications need to demonstrate the following:

- [a] Absence of alternatives. By default, alternative solutions not requiring explicit remoting of link layer indications are preferred, and the burden of proof rests on remoting advocates to show that alternatives (including link-indication aware routing metrics) are unsuitable.

- [b] Conservative behavior. Due to past experience with ICMP "source quench", the burden of proof rests on advocates of remoting proposals to demonstrate that proposals do not violate conservation of packets.

- [c] Security. Remoting proposals need to describe how security issues can be addressed. Where insecure remoted link layer indications are transported over the Internet, an attack can be launched without requiring access to the link.
- [d] Identifiers. When link indications are remoted, it is generally for the purposes of saying something about Internet, Transport or Application layer operations at a remote element. These layers use different identifiers, and so it is necessary to match the link indication with relevant higher layer state. The burden rests on remoting advocates to demonstrate how the link indication can be mapped to the right higher layer state. As an example, if a presence server is receiving remote indications about "Link Up"/"Link Down" status for a particular MAC address, the presence server will need to associate that MAC address with the identity of the user (pres:user@example.com) to whom that link status change is relevant.

## 2.8. Security Considerations

Since link layer indications are typically insecure, proposals incorporating them need to consider the potential security implications of spoofed or modified link layer indications, as well as the potential denial of service attacks. This is particularly important in situations where insecure link layer indications are as a substitute for secure mechanisms operating at a higher layer.

For example, within [[IEEE80211f](#)], "Link Up" is considered to occur when an Access Point sends a Reassociation Response. At that point, the AP sends a frame with the station's source address to a multicast address, thereby causing switches within the Distribution System to learn the station's MAC address, enabling forwarding of frames to the station at the new point of attachment. Unfortunately, this does not take security into account, since the station is not capable of sending and receiving IP packets on the link until completion of the

key exchange protocol defined in [[IEEE80211i](#)]. As a result, link layer indications as implemented in [[IEEE80211f](#)] enable an attacker to disassociate a station located anywhere within the ESS, simply by sending a Reassociation Request frame.

Another example of the potential security implications of link layer indications occurs within DNav4, where link layer indications are used for optimization of IP configuration, rather than using a secured configuration mechanism such as authenticated DHCP [[RFC3118](#)], thereby increasing vulnerability to spoofing.

### [3.](#) Further work

While Figure 1 presents an overview of how link indications are consumed by the Internet, Transport and Application layers, further work is needed to investigate this in more detail.

Given that recent proposals such as [[IEEE80211e](#)] incorporate burst ACKs, the relationship between 802.11 link throughput and frame loss is growing more complex, which may necessitate the development of revised routing metrics, taking the more complex transmission behavior as well as the negotiated rate into account.

At the link and Internet layers, more work is needed to reconcile pre and post-connection metrics, such as reconciling metrics utilized in handoff (e.g. signal strength and link utilization) with link-aware routing metrics (e.g. frame loss and negotiated rate).

At the Transport layer, more work is needed to understand how to react Internet layer indications such as path changes. For example, in an early draft of DCCP [[DCCP](#)], a "Reset Congestion State" option was proposed in [Section 4](#). This option was removed in part because the conditions under which it was to be used were not fully understood:

An HC-Receiver sends the Reset Congestion State option to its sender to force the sender to reset its congestion state -- that is, to "slow start", as if the connection were beginning again. ...

The Reset Congestion State option is reserved for the very few cases

when an endpoint knows that the congestion properties of a path have changed. Currently, this reduces to mobility: a DCCP endpoint on a mobile host MUST send Reset Congestion State to its peer after the mobile host changes address or path.

It may also make sense for the Transport layer to adjust transport parameter estimates in response to "Link Up"/"Link Down" indications and frame loss. For example, it is unclear that the Transport layer should adjust transport parameters as though congestion were detected when loss is occurring in the link layer or a "Link Down" indication has been received.

Finally, more work is needed to determine how link layers may utilize information from the transport layer. For example, it is undesirable for a link layer to retransmit so aggressively that the link layer round-trip time approaches that of the end-to-end transport connection.

## [4.](#) References

### [4.1.](#) Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

### [4.2.](#) Informative 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.

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

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, D. and E. Lear, "Address Allocation for Private Internets", [RFC 1918](#), February 1996.

- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), March 1997.
- [RFC2778] Day, M., Rosenberg, J., Sugano, H., "A Model for Presence and Instant Messaging", [RFC 2778](#), February 2000.
- [RFC2861] Handley, M., Padhye, J. and S. Floyd, "TCP Congestion Window Validation", [RFC 2861](#), June 2000.
- [RFC2914] Floyd, S., "Congestion Control Principles", [RFC 2914](#), [BCP 41](#), September 2000.
- [RFC3118] Droms, R. and B. Arbaugh, "Authentication for DHCP Messages", [RFC 3118](#), June 2001.
- [RFC3428] Campbell, B., Rosenberg, J., Schulzrinne, H., Huitema, C. and D. Gurle, "Session Initiation Protocol (SIP) Extension for Instant Messaging", [RFC 3428](#), December 2002.
- [RFC3748] Aboba, B., , "Extensible Authentication Protocol (EAP)", [RFC 3748](#), June 2004.
- [RFC3775] Johnson, D., Perkins, C. and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.

- [RFC3921] Saint-Andre, P., "Extensible Messaging and Presence protocol (XMPP): Instant Messaging and Presence", [RFC 3921](#), October 2004.
- [RFC3927] Cheshire, S., Aboba, B. and E. Guttman, "Dynamic Configuration of Link-Local IPv4 Addresses", [RFC 3927](#), October 2004.
- [802.11fh] McCann, P., "Mobile IPv6 Fast Handovers for 802.11 Networks", [draft-ietf-mipshop-80211fh-01.txt](#), Internet draft (work in progress), July 2004.
- [Alimian] Alimian, A., "Roaming Interval Measurements", 11-04-0378-00-roaming-intervals-measurements.ppt, IEEE 802.11

submission (work in progress), March 2004.

- [Aguayo] Aguayo, D., Bicket, J., Biswas, S., Judd, G. and R. Morris, "Link-level Measurements from an 802.11b Mesh Network", SIGCOMM '04, September 2004, Portland, Oregon.
- [DCCP] Kohler, E., Handley, M. and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", Internet drafts (work in progress), [draft-ietf-dccp-spec-07.txt](#), July 2004.
- [DNAv4] Aboba, B., "Detection of Network Attachment in IPv4", [draft-ietf-dhc-dna-ipv4-08.txt](#), Internet draft (work in progress), July 2004.
- [EAPIKEv2] Tschofenig, H., D. Kroesenberg and Y. Ohba, "EAP IKEv2 Method", [draft-tschofenig-eap-ikev2-03.txt](#), Internet draft (work in progress), February 2004.
- [Eckhardt] Eckhardt, D. and P. Steenkiste, "Measurement and Analysis of the Error Characteristics of an In-Building Wireless Network", SIGCOMM '96, August 1996, Stanford, CA.
- [Eggert] Eggert, L., Schuetz, S. and S. Schmid, "TCP Extensions for Immediate Retransmissions", [draft-eggert-tcpm-tcp-retransmit-now-00.txt](#), Internet draft (work in progress), July 2004.
- [ETX] Douglas S. J. De Couto, Daniel Aguayo, John Bicket, and Robert Morris, "A High-Throughput Path Metric for Multi-Hop Wireless Routing", Proceedings of the 9th ACM International Conference on Mobile Computing and Networking (MobiCom '03), San Diego, California, September 2003.

- [GenTrig] Gupta, V. and D. Johnston, "A Generalized Model for Link Layer Triggers", submission to IEEE 802.21 (work in progress), March 2004, available at:  
[http://www.ieee802.org/handoff/march04\\_meeting\\_docs/Generalized\\_triggers-02.pdf](http://www.ieee802.org/handoff/march04_meeting_docs/Generalized_triggers-02.pdf)



Institute of Electrical and Electronics Engineers, "Local and Metropolitan Area Networks: Port-Based Network Access Control", IEEE Standard 802.1X, November 2004.

[IEEE80211]

Institute of Electrical and Electronics Engineers, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", IEEE Standard 802.11, 2003.

[IEEE80211e]

Institute of Electrical and Electronics Engineers, "Amendment 7: Medium Access Control (MAC) Quality of Service (QoS) Enhancements", IEEE 802.11e, October 2003.

[IEEE80211f]

Institute of Electrical and Electronics Engineers, "IEEE Trial-Use Recommended Practice for Multi-Vendor Access Point Interoperability via an Inter-Access Point Protocol Across Distribution Systems Supporting IEEE 802.11 Operation", IEEE 802.11f, June 2003.

[IEEE80211i]

Institute of Electrical and Electronics Engineers, "Supplement to Standard for Telecommunications and Information Exchange Between Systems - LAN/MAN Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Specification for Enhanced Security", IEEE 802.11i, November 2004.

[Kim]

Kim, K., Park, Y., Suh, K., and Y. Park, "The BU-trigger method for improving TCP performance over Mobile IPv6", [draft-kim-tsvwg-butrigger-00.txt](#), Internet draft (work in progress), August 2004.

[Kotz]

Kotz, D., Newport, C. and C. Elliot, "The mistaken axioms of wireless-network research", Dartmouth College Computer Science Technical Report TR2003-467, July 2003.

[MIPEAP]

Giaretta, C., Guardini, I., Demaria, E., Bournelle, J., and M. Laurent-Maknavicius, "MIPv6 Authorization and Configuration based on EAP", [draft-giaretta-mip6-authorization-eap-01.txt](#),

Internet draft (work in progress), July 2004.

- [Mishra] Mitra, A., Shin, M., and W. Arbaugh, "An Empirical Analysis of the IEEE 802.11 MAC Layer Handoff Process", CS-TR-4395, University of Maryland Department of Computer Science, September 2002.
- [PEAP] Palekar, A., et al, "Protected EAP Protocol (PEAP) Version 2", [draft-josefsson-pppext-eap-tls-eap-09.txt](#), Internet draft (work in progress), October 2004.
- [Park] Park, S., Njedjou, E. and N. Montavont, "L2 Triggers Optimized Mobile IPv6 Vertical Handover: The 802.11/GPRS Example", [draft-daniel-mip6-optimized-vertical-handover-00.txt](#), July 2004.
- [PPPIANA] Schryver, V., "IANA Considerations for the Point to Point Protocol (PPP)", [draft-schryver-pppext-iana-01.txt](#), August 2003.
- [Shortest] Douglas S. J. De Couto, Daniel Aguayo, Benjamin A. Chambers, and Robert Morris, "Performance of Multihop Wireless Networks: Shortest Path is Not Enough", Proceedings of the First Workshop on Hot Topics in Networking (HotNets-I), Princeton, New Jersey, October 2002.
- [TRIGTRAN] Dawkins, S., et al., "Framework and Requirements for TRIGTRAN", [draft-dawkins-trigtran-framework-00.txt](#), Internet draft (work in progress), August 2003.
- [Vatn] Vatn, J., "An experimental study of IEEE 802.11b handover performance and its effect on voice traffic", TRITA-IMIT-TSLAB R 03:01, KTH Royal Institute of Technology, Stockholm, Sweden, July 2003.
- [Yegin] Yegin, A., "Link-layer Triggers Protocol", [draft-yegin-l2-triggers-00.txt](#), Internet Draft (work in progress), June 2002.
- [Velayos] Velayos, H. and G. Karlsson, "Techniques to Reduce IEEE 802.11b MAC Layer Handover Time", TRITA-IMIT-LCN R 03:02, KTH Royal Institute of Technology, Stockholm, Sweden, April 2003.
- [Vertical] Zhang, Q., Guo, C., Guo, Z. and W. Zhu, "Efficient Mobility Management for Vertical Handoff between WWAN and WLAN", IEEE

INTERNET-DRAFT

Link Layer Indications

14 October 2004

Communications Magazine, November 2003.

[Villamizar]

Villamizar, C., "OSPF Optimized Multipath (OSPF-OMP)", [draft-ietf-ospf-omp-02.txt](#), Internet draft (work in progress), February 1999.

[Appendix A](#). IAB Members at the time of this writing

Bernard Aboba  
Rob Austein  
Leslie Daigle  
Patrik Falstrom  
Sally Floyd  
Mark Handley  
Bob Hinden  
Geoff Huston  
Jun-Ichiro Itojun Hagino  
Eric Rescorla  
Pete Resnick  
Jonathan Rosenberg

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 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.

## Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS

IAB

Informational

[Page 31]

---

INTERNET-DRAFT

Link Layer Indications

14 October 2004

OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY 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.

## Copyright Statement

Copyright (C) The Internet Society (2004). 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.

