Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: 20 September 2026

M. Smith

19 March 2026

# IPv6 Over Nothing
## draft-smith-6man-ipv6-over-nothing-01

## Abstract

A perspective on the function of the network layer is that it abstracts away the differences between the various underlying link layer frame and addressing formats, unifying them into a common protocol data unit format and addressing scheme, namely the IPv4 or IPv6 protocols, and hiding those details from the upper transport layer protocols. As IPv6 is expected to become the dominant network layer protocol, and Ethernet has become the dominant link layer protocol, this memo proposes eliminating the overhead of the abstraction of Ethernet by IPv6 and using IPv6 directly as both the link layer and network layer protocol.

## Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 20 September 2026.

## Copyright Notice

Table of Contents



## 1. Introduction

> "In protocol design, perfection has been reached not when there is nothing left to add, but when there is nothing left to take away." [RFC1925]

As observed in the Internet Protocol Suite overview in [RFC6272], "The Internet layer provides a uniform network abstraction network that hides the differences between various network technologies."

More specifically, the Internet layer is abstracting away the different frame formats, processing and addressing formats of the various link layers that are used to construct an IP based network. To layers above the network layer, there is a uniform packet format, uniform processing, and a uniform addressing scheme that identifies all packet sources and destinations.

Ethernet [IEEE802.3] has become the most widely deployed physical link layer technology in IP networks, displacing other link layer technologies such as ATM, SONET/SDH/PoS, Token Ring and Frame Relay.

This is due to Ethernet becoming the most popular link layer technology used to deploy wired Local Area Networks used to interconnect commodity desktop personal computers. Ethernet became so commodified that it became both the cheapest way to interconnect IP devices, including routers within IP networks, and also became the best link layer technology to enhance with performance improvements, since it provided the largest likely market for those improvements.

A useful computer science principle is to optimise for the common case. Since IPv6 is expected to become the common case IP protocol in use, Ethernet is the common case for interconnecting network devices, and the IP layer's purpose is to abstract away different link layer's differences, there is an opportunity to optimise for the common case of IPv6 over Ethernet and use IPv6 directly as a link layer protocol. This would eliminate the need to abstract away Ethernet's link layer frame and addressing properties, simplifying the network and increasing performance. It would also eliminate the need to also run a link layer redundancy protocol such as the Spanning Tree Protocol [STP].

This memo describes how this can be achieved.

Note that this is a serious proposal, despite it sounding like an April Fools memo. Should it progress to being published as an RFC, it could be published as Standards Track RFC on April Fools for the amusement of the audience.

### 1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

## 2. Terminology

IPv6 link - a link layer link implementing this specification.

Maximum Transmission Unit (MTU) - the maximum number of octets available to the next protocol layer above to carry its data. Could also be described as the "Payload MTU" for the current protocol layer.

## 3. Link Type

When using IPv6 as a link layer protocol, an IPv6 link is a point-to-point link, per [RFC4861]’s point-to-point link definition. There is no such thing as a multi-access link connecting more than 2 nodes when using IPv6 as a link layer protocol.

Also as per [RFC4861], these point-to-point links are multicast capable, by sending IPv6 multicasts to the far end of the point-to-point link, and optionally looping the multicast back to the sending host itself when necessary [Reference to IPv6 version of RFC1112].

If an IPv6 link is used between an host and a router, this means that an IPv6 host is directly connected to a single upstream IPv6 router. If a host requires network attachment redundancy, it will need to be connected to two different routers via two different point-to-point IPv6 links.

## 4. IPv6 Addressing

As per [RFC4291], all interfaces attached to IPv6 links are required to have at least one Link-Local unicast address. These addresses are typically generated and configured via SLAAC [RFC4862][RFC8064].

Link-local addresses are the analogue for traditional link layer addresses such as Ethernet MAC addresses, as packets with link-local addresses are limited to a single link [RFC4007].

An interface attached to an IPv6 link may and likely will gain other GUA [RFC4291] and/or ULA [RFC4193] addresses, and perhaps additional link-local addresses, via IPv6 address configuration methods such as SLAAC [RFC4862][RFC8064], DHCPv6 [RFC9915] or manual configuration.

## 5. IPv6 Link Error Detection

### 5.1. Inter-Router Link Error Detection

To detect packet loss and link corruption that causes packet loss, isolated to a single link between two routers, single hop IPv6 Bidirectional Forwarding Detection (BFD) is used [RFC5881]. Asynchronous mode MUST be used.

BFD packets are carried in UDP [RFC768], with the UDP checksum protecting the BFD protocol in the UDP payload.

Per [RFC3819], link layer checksums are much more robust than the 16-bit 1’s compliment checksum used by UDP, TCP, etc.

As BFD is being used as the link error detection mechanism over an IPv6 link, as a proxy for traditional and much stronger link layer checksums [RFC3819], the UDP Additional Payload Checksum (APC) option [RFC9868] MUST be used for BFD Control packets when available on both routers.

Alternatively, BFD Authentication [RFC5880] MUST be used if the UDP APC option is not available on both routers. The MD5 or SHA1 hash functions MUST be used to authenticate BFD packets, as they will also detect link corruption that causes packet loss. Simple Password authentication MUST NOT be used, as it will not detect link errors.

If the BFD Echo function is used, then the UDP Additional Payload Checksum (APC) option [RFC9868] SHOULD be used, to more quickly detect link failure (that was not detected via the UDP checksum), rather than it being detected via the next Asynchronous mode periodic BFD Control packet protected via the UDP APC Option or BFD Authentication.

Note that the IPsec Authentication Header in Transport Mode [RFC4302] could also be used to protect BFD Echo packets should the UDP APC option not be available, with the BFD Echo packet sender authenticating its own looped back packets.

### 5.2. Host-Router and Router-Host Link Error Detection

Transport layer protocols such as TCP [RFC793] and UDP [RFC768] use an end-to-end checksum to detect errors that have occurred while packets travel over the network between the transport layer end-points residing in the source and destination hosts.

As inter-router IPv6 links are protected via BFD (see previous section), or link layer checksums for other link layer link types, if no inter-router links are suffering from errors, then it can be deduced that a fault lies with either the link between a source host and its upstream router, or the final router and its downstream destination host, or both of these links. Conventional ICMPv6 Echo Requests and Replies ("ping") [RFC4443] can be used to determine which of these two links is faulty or in fact if they both are.

If necessary, BFD [RFC5881] could be used between the hosts and their routers to better check for link errors, following the same implementation requirements as when BFD is being used between two routers, as discussed previously.

## 6. Maximum Transmission Unit

The (Payload) Maximum Transmission Unit for an IPv6 link is variable, up to the maximum payload size that can be carried in an IPv6 packet.

The default Maximum Transmission Unit for an IPv6 link is 2^16-1 or 65535 octets, the maximum value for the IPv6 Payload Length field [RFC8200].

The maximum Maximum Transmission Unit for an IPv6 link is 2^32-1 or 4294967295 octets, the maximum value for the Jumbo Payload Length field when using the Jumbo Payload IPv6 Hop-by-Hop option [RFC2675].

The total IPv6 packet size sent between nodes will be the value of either the IPv6 Payload Length field or the Jumbo Payload Length field plus 40 octets, the size of the IPv6 Header [RFC8200] (note that the Payload Length or Jumbo Payload Length field values include the length of any optional Extension Headers behind the IPv6 Header).

## 7. Neighbor Discovery

Although IPv6 links do not have actual link layer addresses, and are point-to-point, [RFC4861] Neighbor Discovery still needs to be performed, as specified for point-to-point links in [RFC4861]. This is because Neighbor Discovery, in addition to normally resolving an IPv6 address into a link layer address for link layers with addresses, also discovers whether or not an IPv6 address exists at the far end of the IPv6 (point-to-point) link, and monitors the continued existence of discovered IPv6 addresses via Neighbor Unreachability Discovery (NUD) [RFC4861].

Neighbor Discovery is not optional because it prevents the so-called "ping-pong" problem described in [RFC6164] on inter-router links. This problem exists because router IPv6 implementations weren’t or aren’t performing Neighbor Discovery on point-to-point links. Instead, they were or are assuming that any and all non-local addresses must exist at the far end of the point-to-point link. In other words, they weren’t testing for the existence of an IPv6 address via Neighbor Discovery at the far end of the link before sending packets to that IPv6 address.

As IPv6 addresses are carried directly within the Neighbor Discovery protocol, and they are the functional equivalent of link layer addresses, the Source/Target Link-Layer Address [RFC4861] option is not required in either Neighbor Solicitations or Neighbor Advertisements.

## 8. Router Discovery

As could be expected from the previous Neighbor Discovery section, [RFC4861] router discovery works as usual, with the exception that Router Solicitations and Router Advertisments do not contain Source/ Target Link-layer Address options.

## 9. Carrying Other Link Layer and Network Layer Protocols

It may be useful or necessary to carry link layer or network layer protocols over an IPv6 link. For example, it may be useful to use Link Layer Discovery Protocol [LLDP] between routers attached to an IPv6 link, or transport legacy IPv4 over IPv6 links.

Carrying these protocols over IPv6 links can be achieved via various "tunnelling" over IPv6 methods such as [RFC2473]. For example, LLDP can be carried in Ethernet frames that are tunnelled over IPv6, with a suitable IPv6 unicast or multicast destination address.

Note that as IPv6 supports multiple different multicast scopes that cover increasingly larger domains of the IPv6 network [IANA-IPv6-MCAST-SCOPES], it would be possible to tunnel multicast link layer protocols such as LLDP beyond a single link, unlike when they are carried in traditional link layer protocols such as Ethernet. In this case, multicast routing would need to be enabled on the network to send these types of multicast IPv6 packets beyond a single IPv6 link.

## 10. IPv6 Over Almost Nothing

It is common for a networking stack to fill in the Destination Address, Source Address and Ethernet Type fields of each frame before transmission. Upon frame reception, a Ethernet Network Interface Card (NIC) uses a Destination Address matching mechanism within the NIC to determine if it should accept the frame and pass it onto the networking protocol stack for further processing.

If an Ethernet NIC is put into "promiscuous" mode, it will accept frames with any Destination Address value, passing them onto the networking protocol stack.

Over a point-to-point Ethernet link, once a NIC is in promiscuous mode, the Destination and Source Address fields of an Ethernet frame could be used to carry information other than Ethernet addresses. This would result in an additional 12 octets per frame being available.

Furthermore, again over a point-to-point Ethernet link, the 2 octet Ethernet type field could also be used to carry information other than Ethernet frame type information.

These initial 14 octets in the Ethernet frame could be used to carry the first 14 octets of the IPv6 packet, making this Ethernet link an "IPv6 over Almost Nothing" link, a variant of the IPv6 link described previously.

Note that Ethernet’s maximum (payload) MTU limitations would still apply, although the MTU available would be 14 octets larger. The Ethernet Frame Check Sequence would also still be present in the frames, meaning that BFD would not be necessary on a link to detect packet loss and corruption that causes packet loss.

It is important to ensure that the point-to-point link is a true physical link, or if not, there are no devices present within the link that will interpret the values of the Destination and Source Addresses and the Ethernet type field in these IPv6 over almost nothing frames. If such a device was present it will likely be confused by the IPv6 header data now contained where the former Ethernet Destination and Source address and Type information resided, and possibly drop rather than forward these frames. The device may even get so confused it ceases to operate.

## 11. IANA Considerations

This memo includes no request to IANA.

## 12. Security Considerations

This document should not affect the security of the Internet.

## 13. References

### 13.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

### 13.2. Informative References

Acknowledgements

Thanks to the following people for their reviews and comments (in alphabetical order): Mark Baylis, Jonathan Brewer, Tom Hill, Adam Kalisz, Patrik "Jima" Laughton, Dave Seddon, Blake Willis.

Author’s Address

Mark Smith
PO BOX 521
Heidelberg Victoria 3084
Australia
Email: markzzzsmith@gmail.com