

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 20 September 2026

F. L. Templin, Ed.
The Boeing Company
19 March 2026

AERO/OMNI Base Specification Amendments (Volume 1)
draft-templin-6man-aero-omni-amen-03

Abstract

The Automatic Extended Route Optimization (AERO) and Overlay Multilink Network (OMNI) Interface functional specifications have reached a level of maturity ready for advancement in the RFC publication process. Updates to the base specifications are documented in this first amendment and any additional future amendments as necessary.

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

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. AERO/OMNI: First Amendment	2
2.1. Amendment 1.1: Prefix Delegation Administration	3
2.2. Amendment 1.2: AERO/OMNI MLA Addressing	4
2.3. Amendment 1.3: Off-link Nature of Non-MLAs	5
2.4. Amendment 1.4: OMNI Interface Virtual Router	6
2.5. Amendment 1.5: OMNI Interface Neighbor Cache	6
2.6. Amendment 1.6: Scalable Mapping	6
2.7. Amendment 1.7: Address Duplication Implications	8
2.8. Amendment 1.8: Setting Auth Offset	8
2.9. Amendment 1.9: Surrogate Multilink Pilots (MPs)	8
2.10. Amendment 1.10: Extended AFVI TLVs	9
2.11. Amendment 1.11: Path Change Mitigations	10
2.12. Amendment 1.12: MAP Proxy/Servers	11
3. IANA Considerations	12
4. Security Considerations	12
5. Acknowledgements	12
6. References	12
6.1. Normative References	12
6.2. Informative References	13
Appendix A. Change Log	14
Author's Address	14

1. Introduction

The Automatic Extended Route Optimization (AERO) [I-D.templin-6man-aero3] and Overlay Multilink Network (OMNI) Interface [I-D.templin-6man-omni3] functional specifications have reached a level of maturity ready for advancement in the RFC publication process, while this and any other future amendments provide normative appendices that update the AERO/OMNI base specifications.

These amendments assume consistency with IPv4 [RFC0791], IPv6 [RFC8200] and the IPv6 neighbor discovery and autoconfiguration services [RFC4861][RFC4862][RFC9915]. The term "IP" refers universally to either the IPv4 or IPv6 protocol version.

2. AERO/OMNI: First Amendment

Amendments to the AERO/OMNI specifications appear in the following sections in the chronological order in which they were identified, formalized and recorded. Additional amendments may appear in future volumes.

2.1. Amendment 1.1: Prefix Delegation Administration

The AERO/OMNI base specifications include comprehensive instructions for Clients to request and receive Mobile Network Prefix (MNP) IP Prefix Delegations (PDs) from the mobility service. The specifications suggest that the Client should apply these MNP PDs on downstream-attached (or, "tethered") End User Network (EUN) interfaces but do not discuss specific MNP administrative procedures.

Per this amendment, when a Client receives an MNP PD it should provision the MNP over EUNs in a manner consistent with [RFC9663] and [RFC9762]. More specifically, the Client assigns some portions of the MNP to its EUN interfaces and optionally also sub-delegates other portions of the MNP to requesting EUN nodes. The Client can also use virtual interfaces (such as a loopback) as EUN interfaces. The Client should then assign an IP address taken from each EUN prefix to the corresponding EUN interface using standard IP address (auto)configuration procedures.

The Client then associates the IPv6 Subnet Router Anycast (SRA) address [RFC4291] corresponding to the MNP with the OMNI interface but does not assign it to the interface. For example, if the Client receives the IPv6 MNP PD 2001:db8:1:::/48, it can either accept packets with an SRA address 2001:db8:1:*::/64 as the destination and respond with packets that use one of the Client's EUN addresses as the source or forward the packet to an EUN router that configures a more-specific sub-prefix of the ::/48. In the first case, the wildcard response would represent all sub-prefixes as reachable when in fact some may not be assigned to any EUN links; in the latter case, an adversary may be able to map the EUN internal subnets.

After the Client configures the SRA address and assigns EUN addresses it can operate as an IP host over the OMNI interface according to the weak end system model [RFC1122] while also serving as an IP router for its EUNs. The Client then engages host and router operations the same as per [RFC4861] and [RFC8200] except that the Client engages the OMNI interface as a combined host/router interface.

When the Client acts as a host over the OMNI interface, it can send Router Solicitations (RSs) to elicit Router Advertisements (RAs) from OMNI link Proxy/Servers. The Client can then use its EUN addresses for packets exchanged between its local applications and correspondents reached via OMNI link neighbors. (When the Client has no EUN addresses, it can instead use its OMNI interface Multilink Local Address (MLA) [I-D.templin-6man-mla] with the understanding that the packets may be routable only within the OMNI link limited domain.)

When the Client acts as a router over the OMNI interface, it forwards IP packets between EUN peers and correspondents reached via OMNI link neighbors but never sends RA messages over the OMNI interface. This may require the Client to enable IP forwarding on the OMNI interface but without representing itself as a router in OMNI link IPv6 Neighbor Discovery (IPv6 ND) messages. The Client instead sends RAs over its EUN interfaces that include EUN portions of the MNP in Prefix Information Options (PIOs) and also represents itself as a router in other EUN interface IPv6 ND messages.

Note that the Client could also optionally assign each EUN address it configures to the OMNI interface. This would give the outward appearance of strong end system support [RFC1122], albeit with added complexity and ambiguity for the Client to coordinate the same IP unicast address assigned to multiple interfaces.

Further details on MNP PD administrative options are beyond the scope of this amendment.

2.2. Amendment 1.2: AERO/OMNI MLA Addressing

In the addressing model supported under the current specifications, the MLAs of all nodes within the OMNI link limited domain (Clients, Proxy/Servers and others) should be reachable by all other nodes on the link. MLAs of nodes beyond the extent of the OMNI link will appear as unreachable destinations though they may be reachable within other OMNI links. This is due to the limited domain nature of OMNI link routing for MLAs and highlights why MLAs must appear as a lower preference than other GUAs in address selection policies.

In this addressing model, an MLA prefix (e.g., 2001:30::/28 [RFC9374]) is configured on-link on the OMNI interface and any MLA routes discovered by Mobile Ad-hoc NETworking (MANET) routing protocols are maintained in an alternate routing table instead of the main kernel routing table. The MANET routing protocol can then manage its multihop routes dynamically as necessary while the network layer forwards all packets with MLA destinations into the OMNI interface. The on-link nature of the MLA prefix will cause the network layer to invoke address resolution for destination MLAs within the limited domain. This will cause the OMNI Address Resolution Source (ARS) to propagate NS(AR) messages over the OMNI link to generate an NA(AR) response from an Address Resolution Target (ART).

The ARS's NS(AR) uses the invoking packet's address or its local MLA as the Source, uses the MLA of the target as the Target and uses the solicited-node multicast address as the Destination. The ART's NA(AR) uses the NS(AR) target as both the Source and Target and uses

the NS(AR) source as the Destination. AERO routing will cause these messages to traverse the OMNI link as encapsulated OAL packets and return fresh address resolution information.

2.3. Amendment 1.3: Off-link Nature of Non-MLAs

Non-MLAs include MNP/FNP addresses and other GUA addresses matched only by "default". These prefixes must be configured as reachable (but not on-link) via the OMNI link. Clients therefore configure routes in the IPv6 routing system that cover MNPs and/or FNPs and with next hop set to the Link-Local Address (LLA) of the OMNI interface internal virtual router. This will cause all packets with destinations within these off-link prefixes to be delivered to the virtual router. The route(s) may include specific MNP/FNP prefixes or could even be `::/0` (i.e., default) and/or the full Mobility Service Prefix.

The OMNI link virtual router then acts as an ARS to resolve adaptation layer addressing information for the packet's destination prefix internally within the adaptation layer and without disturbing the network layer. While address resolution is in progress, the virtual router maintains a short queue (possibly only a single entry) of packets destined to the ART prefix the same as for [RFC4861]. The virtual router should not attach subject packets awaiting address resolution to NS(AR) messages, as they may be produced by high-volume applications (e.g., "flood-pings") that send many unacknowledged packets without waiting for a response.

The ARS's NS(AR) message uses the subject packet's source address as the Source, the subject packet's destination address as the Target and the solicited-node multicast address as the Destination. The ART's responsive NA(AR) message uses its own MLA as the Source, the Subnet Router Anycast (SRA) address that matches the NS(AR) target as the Target and the NS(AR) source as the Destination.

The ART therefore responds to all addresses within the IPv6 prefix that covers the NS(AR) Target Address while including a Route Information Option (RIO) which asserts reachability for an entire IPv6 prefix whereas the target itself (or even the target's more-specific prefix) may be unreachable.

This non-MLA address resolution allows the ARS to discover information for entire IPv6 prefix ranges without necessarily conveying reachability information for specific destinations (or even sub-prefixes) within the prefix. The system therefore depends on the target network returning ICMPv6 Destination Unreachable messages for unreachable destinations within (reachable) prefixes.

2.4. Amendment 1.4: OMNI Interface Virtual Router

The virtual router entity within the OMNI interface must present itself to the network layer as a minimally qualified IPv6 router according to IPv6 node requirements [RFC8504]. This includes using its internal LLA to send solicited and unsolicited Router Advertisements as well as respond to Neighbor Solicitations by returning solicited Neighbor Advertisements.

The virtual router entity should also provide a configuration option allowing it to either respond to or ignore ICMPv6 Echo Request messages addressed to its internal LLA or the subnet router anycast address for its MLA prefix(es).

The virtual router entity uses its internal LLA as the Source or Destination Address for message exchanges with the network layer. The virtual router entity need not respond to messages addressed to fe80::/128 as this represents an ambiguous LLA form not widely acknowledged as a valid IPv6 unicast/anycast address.

2.5. Amendment 1.5: OMNI Interface Neighbor Cache

Each OMNI interface maintains a standard IPv6 network layer neighbor cache the same as for any IPv6 interface and also maintains an adaptation layer neighbor cache internally. The network layer neighbor cache maintains entries (NCEs) for the adaptation layer as a virtual router as well as for active on-link destinations only, while the adaptation layer neighbor cache maintains NCEs for both on- and off-link destinations reached via the OMNI interface.

Each network layer NCE resolves to the singular OMNI interface internal link-layer address; this means that all NCE destinations would appear to belong to the same (singular) neighbor. The adaptation layer virtual router entity will then map the network layer NCE to the corresponding adaptation layer NCE by examining the IP destination address rather than the link-layer address. This relationship establishes a 1x1 mapping between the network layer as a virtual host and the adaptation layer as a virtual peer host/router on a shared link.

2.6. Amendment 1.6: Scalable Mapping

Scaling properties for the worldwide civil aviation airplane population are likely to remain within reasonable bounds for the pure BGP routing system discussed in [I-D.templin-6man-aero3] for the foreseeable future. However, the advent of unmanned air systems and all other manners of mobile nodes will soon present multiple orders of magnitude more mobility targets which may exceed the carrying

capacity of a BGP-only service.

In order to support unbounded scaling, the BGP routing system can be limited to carry only the MLAs of all Proxy/Servers on the OMNI link (and possibly also the MNPs/FNPs/MLAs of a limited number of mobile nodes) without carrying the entire population of mobile node MNP/FNP/MLA information. Each MAP Proxy/Server then registers the MNP/FNP routes and MLA addresses of its dependent mobile nodes with a scalable mapping system that can be used to resolve a target address based on longest prefix match into a MAP Proxy/Server MLA address. The Domain Name System (DNS) ip6.arpa reverse zone can be used for this purpose as suggested in [RFC9374].

Address resolution then becomes a two-phase operation where the address resolution request is first forwarded (e.g., via a default or more-specific route) to a Gateway which consults the DNS to determine the MLA of the current MAP Proxy/Server for the address resolution target. The mapping system agent then changes the OAL destination to the discovered MLA and forwards the address resolution request to the MAP which returns a fully qualified address resolution response.

By maintaining mobile node to MAP mappings in a scalable ancillary lookup directory database, the BGP routing system only needs to scale to the total population of Proxy/Servers and Gateways that make up the OMNI link (plus possibly also a limited number of mobile nodes). This is likely to remain within acceptable scaling limits even for extremely large mobile node populations for the foreseeable future.

Note that the Gateway performs these mapping system lookups only for subject prefixes associated with the OMNI link, e.g., those covered by MSPs or any well-known FNPs. For other subject prefixes that match only "default" the Gateway returns an address resolution response with a Route Information Option (RIO) that includes a /64 prefix covering the target address while including its own MLA as the address resolution result.

Note also that MAP Proxy/Servers need only add or replace DNS resource records with the IP prefix mappings as they receive registration requests from new Clients. If the Client moves to a new MAP Proxy/Server, the new MAP simply replaces the old resource records with fresh information. If the resource records expire before a new MAP Proxy/Server supplies fresh information, the records are removed.

Global propagation delays for DNS resource record updates and the location privacy considerations for mobile node affiliations with MAP Proxy/Servers suggest that the OMNI link should configure a "two-faced" DNS service infrastructure maintained separately from the

global public DNS. Such a service should be optimized for fast updates so that Gateways always receive fresh address resolution information.

2.7. Amendment 1.7: Address Duplication Implications

An implicit assumption in AERO/OMNI is that IP address duplication within the same OMNI link domain could result in an unresolvable and harmful interaction for any nodes affected. The AERO/OMNI routing system and neighbor discovery services are founded on an expectation of uniqueness of the MNPs/FNPs/MLAs claimed by nodes. The MLA in particular is carefully coordinated with a registration authority and bound to a public key identity to ensure uniqueness. MNP/FNP delegations are similarly bound to an MLA when delegated by AERO/OMNI infrastructure supporting nodes to prevent duplication.

On the other hand, link-layer address duplication (while presumably rare) would result in simply minor performance congestion since the IP stacks of the nodes with duplicate link-layer addresses would reject any mis-directed packets while the stacks of routers would be able to discern control messages based on their IPv6 MLA addresses as a unique node identifier. While link-layer address duplication is not desirable, it should have no harmful operational effects.

2.8. Amendment 1.8: Setting Auth Offset

The OMNI option trailer includes a 1-octet Auth Offset field that encodes the offset from the beginning of the OMNI sub-options to the first authentication (or "authentication helper") sub-option. This allows for more efficient authentication verification at OAL destinations since they need not wade through potentially long concatenations of leading non-authentication sub-options.

When an OMNI option does not assert an authentication/helper sub-option offset, the OAL source can set Auth Offset to any value no smaller than OMNI Length. For example, the OAL source can simply set the value 'ff' as an unambiguous indication that no offset is asserted.

2.9. Amendment 1.9: Surrogate Multilink Pilots (MPs)

The AERO/OMNI base specifications require that control messages (including IPv6 ND and Multilink Pilot (MP) messages) must be transmitted as whole packets not subject to OAL fragmentation. Very large original IP packets may therefore be unsuited to serve as MP messages.

In that case, the OAL source can continue to hold the large packet in a queue while sending unsolicited Neighbor Advertisement (uNA) messages that include Source, Destination and Flow Label information copied from the large packet as surrogate MPs.

Transmission of surrogate MPs will cause OAL intermediate systems to securely establish AFVI state, after which the OAL source can release the large packet (while applying fragmentation if necessary followed by header compression) to follow the newly-established AFVI path in the data plane.

2.10. Amendment 1.10: Extended AFVI TLVs

The OMNI specification defines a segment Routing Header (SRH) TLV termed the AERO Flow Vector Index (AFVI) TLV. This amendment defines an extended form of the AFVI TLV to include a Sequence Number and Window size as shown below:

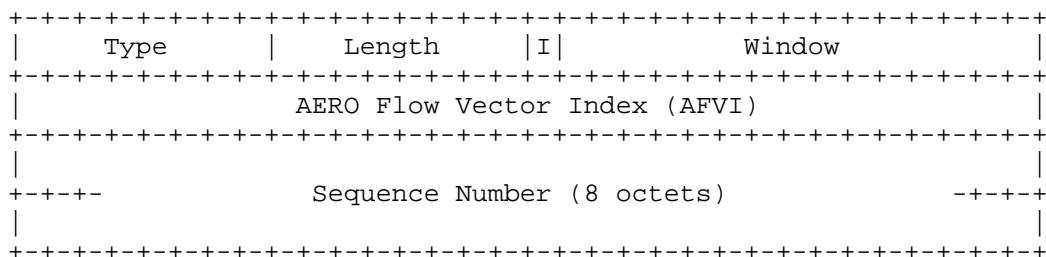


Figure 1: Extended AFVI TLV

In this extended form:

- * Type is set the same as in the OMNI specification.
- * Length is set to 14.
- * I is an "Initialize" flag used for the same purpose as in the OMNI specification.
- * Window is a 15-bit integer that encodes the size of the current send window in units of 512 Sequence Numbers. For example, if Window encodes the value 128 the current send window is $(128 * 512) = 65536$ Sequence Numbers.
- * AERO Flow Vector Index (AFVI) is a 32-bit field used for the same purpose as in the OMNI specification.

- * Sequence Number is an 8-octet field that encodes the same value that would appear in the Extended Fragment Header (EFH). For unfragmented control messages, this value can therefore include a Sequence Number used to establish AFV state.

When the OAL source includes this extended form of the SRH AFVI TLV in unfragmented control messages, it should not also include an Extended Fragment Header (EFH) as it would only include redundant or conflicting information that could confuse OAL intermediate systems. The SRH AFVI TLV also has the distinct advantage that it is covered by the authentication signature included in the SRH HMAC TLV.

This amendment therefore deprecates inclusion of the EFH in unfragmented control messages and mandates inclusion of the SRH extended AFVI TLV in its place.

This amendment further deprecates inclusion of the OMNI Neighbor Synchronization sub-option, as all window synchronization will be unidirectional based on the SRH AFVI TLV and therefore no TCP-like bidirectional handshaking is necessary.

2.11. Amendment 1.11: Path Change Mitigations

The AERO specification includes path change mitigations that permit an OAL intermediate node to revert to sending packets with uncompressed headers when the next hop in the AFV path has changed. The OAL destination is then responsible for returning ICMPv6 Parameter Problem messages with code "Compressed header expected".

In open Internetworks not protected by lower layer security, however, this arrangement could open a Denial of Service (DoS) vector in which an unauthorized source produces a flood of packets with uncompressed headers causing an OAL destination to return false path change reports.

The OAL intermediate node should therefore return an ICMP Parameter Problem message (subject to rate limiting) with code "Path change" (TBD) and then forward packets with uncompressed headers only if the underlay network is secured against DoS spoofing. In open Internetworks, the OAL intermediate node should instead simply drop the packets.

2.12. Amendment 1.12: MAP Proxy/Servers

The OMNI specification implies that all First Hop Segment (FHS) Proxy/Servers should be eligible to also act as distributed Mobility Anchor Points (MAPs). For FHS Proxy/Servers positioned in highly dynamic environments, however, this might lead to undesirable interactions with the OMNI link routing system. These dynamic FHS Proxy/Servers should therefore not accept the MAP role themselves, but should instead act as transparent proxies to forward Client registration requests to MAPs located in more stable environments.

The Client discovers the list of MAPs for the OMNI link by querying the Domain Name System (DNS) for the OMNI link Potential Router List (PRL). The PRL query returns a list of MAP Proxy/Server resource records that include the MLA, underlay IP address(es) and geographic positioning information for each MAP. Clients can then select specific MAPs by placing the MLA in an OMNI Router Solicitation (RS) Destination Address; the FHS Proxy/Server for the Client's link will in turn transparently proxy and forward the RS to the MAP.

Clients may selectively test connectivity to candidate MAPs before selecting one with the desired performance profile. The MAP should in turn defer its interactions with the OMNI link routing system until the Client indicates its intention to commit. This amendment therefore adds a new (C)ommit flag to the OMNI Proxy/Server Control sub-option as shown in Figure 2:

```

+-----+
| Sub-Type=15 | Sub-Length |M|P|N|A|R|C|   Reserved   |
+-----+
~

```

Figure 2: Proxy/Server Control With (C)ommit Flag

When the Client sends an RS to test a candidate MAP Proxy/Server's reachability and performance, it includes a Proxy/Server Control sub-option with the (C)ommit flag set to 0. The MAP Proxy/Server then returns a Router Advertisement (RA) message without committing to serve the Client. When the Client instead sets the (C)ommit flag to 1, the MAP Proxy/Server returns an RA with zero lifetimes if it is unable to commit. Otherwise, the MAP Proxy/Server creates a Neighbor Cache entry, delegates the Client's requested MNP(s) and injects the Client's MNP(s) and MLA into the OMNI link routing system. The MAP Proxy/Server then returns an RA with valid lifetimes.

3. IANA Considerations

This document includes no actions for IANA.

4. Security Considerations

The security considerations in the normative references apply.

5. Acknowledgements

This work is aligned with the Boeing/Virginia Tech National Security Institute (VTNSI) 5G MANET research program.

Honoring life, liberty and the pursuit of happiness.

6. References

6.1. Normative References

[I-D.templin-6man-aero3]

Templin, F., "Automatic Extended Route Optimization (AERO)", Work in Progress, Internet-Draft, draft-templin-6man-aero3-60, 28 February 2026, <<https://datatracker.ietf.org/doc/html/draft-templin-6man-aero3-60>>.

[I-D.templin-6man-mla]

Templin, F., "IPv6 Addresses for Ad Hoc Networks", Work in Progress, Internet-Draft, draft-templin-6man-mla-32, 16 February 2026, <<https://datatracker.ietf.org/doc/html/draft-templin-6man-mla-32>>.

[I-D.templin-6man-omni3]

Templin, F., "Transmission of IP Packets over Overlay Multilink Network (OMNI) Interfaces", Work in Progress, Internet-Draft, draft-templin-6man-omni3-80, 28 February 2026, <<https://datatracker.ietf.org/doc/html/draft-templin-6man-omni3-80>>.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.

- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [RFC9915] Mrugalski, T., Volz, B., Richardson, M., Jiang, S., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", STD 102, RFC 9915, DOI 10.17487/RFC9915, January 2026, <<https://www.rfc-editor.org/info/rfc9915>>.

6.2. Informative References

- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <<https://www.rfc-editor.org/info/rfc1122>>.
- [RFC8504] Chown, T., Loughney, J., and T. Winters, "IPv6 Node Requirements", BCP 220, RFC 8504, DOI 10.17487/RFC8504, January 2019, <<https://www.rfc-editor.org/info/rfc8504>>.
- [RFC9374] Moskowitz, R., Card, S., Wiethuechter, A., and A. Gurtov, "DRIP Entity Tag (DET) for Unmanned Aircraft System Remote ID (UAS RID)", RFC 9374, DOI 10.17487/RFC9374, March 2023, <<https://www.rfc-editor.org/info/rfc9374>>.
- [RFC9663] Colitti, L., Linkova, J., Ed., and X. Ma, Ed., "Using DHCPv6 Prefix Delegation (DHCPv6-PD) to Allocate Unique IPv6 Prefixes per Client in Large Broadcast Networks", RFC 9663, DOI 10.17487/RFC9663, October 2024, <<https://www.rfc-editor.org/info/rfc9663>>.
- [RFC9762] Colitti, L., Linkova, J., Ma, X., Ed., and D. Lamparter, "Using Router Advertisements to Signal the Availability of DHCPv6 Prefix Delegation to Clients", RFC 9762, DOI 10.17487/RFC9762, June 2025, <<https://www.rfc-editor.org/info/rfc9762>>.

Appendix A. Change Log

<< RFC Editor - remove prior to publication >>

Differences from -00 to -01:

- * Added amendments 1.2 through 1.11.

Differences from earlier versions:

- * First draft publication.

Author's Address

Fred L. Templin (editor)
The Boeing Company
P.O. Box 3707
Seattle, WA 98124
United States of America
Email: fltemplin@acm.org