

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 15 December 2025

F. L. Templin, Ed.
The Boeing Company
13 June 2025

Transmission of IP Packets over Overlay Multilink Network (OMNI)
Interfaces
draft-templin-6man-omni3-61

Abstract

Air/land/sea/space mobile nodes (e.g., aircraft of various configurations, terrestrial vehicles, seagoing vessels, space systems, enterprise wireless devices, pedestrians with cell phones, etc.) communicate with networked correspondents over wireless and/or wired-line data links and configure mobile routers to connect end user networks. This document presents a multilink virtual interface specification that enables mobile nodes to coordinate with a network-based mobility service, fixed node correspondents and/or other mobile node peers. The virtual interface provides an adaptation layer service suited for both mobile and more static environments such as enterprise and home networks. Both Provider-Aggregated (PA) and Provider-Independent (PI) addressing services are supported. This document specifies the transmission of IP packets over Overlay Multilink Network (OMNI) Interfaces.

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 15 December 2025.

Copyright Notice

Copyright (c) 2025 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	4
2. Terminology	7
3. Requirements	18
4. Overlay Multilink Network (OMNI) Interface Model	19
5. OMNI Interface Maximum Transmission Unit (MTU)	25
6. The OMNI Adaptation Layer (OAL)	26
6.1. OAL Source Encapsulation and Fragmentation	27
6.2. OAL L2 Encapsulation and Re-Encapsulation	31
6.3. Reassembly and Decapsulation	34
6.4. OMNI-Encoded IPv6 Extension Headers	36
6.5. OMNI Full and Compressed Headers	39
6.6. L2 UDP/IP Encapsulation Avoidance	44
6.7. OAL Identification Window Maintenance	44
6.8. OAL Fragmentation Reports and Retransmissions	50
6.9. OMNI Interface MTU Feedback Messaging	51
6.10. OAL Composite Packets	53
6.11. OAL Bubbles	55
6.12. OAL Requirements	55
6.13. OAL Fragmentation Security Implications	56
6.14. Control/Data Plane Considerations	57
7. Ethernet-Compatible Link Layer Frame Format	58
8. OMNI Addressing	59
9. Node Identification	63
10. Address Mapping - Unicast	63
10.1. The OMNI Option	65
10.2. OMNI Sub-Options	68
10.2.1. Pad1	69
10.2.2. PadN	70
10.2.3. Node Identification	70
10.2.4. CGA	72
10.2.5. RSA Signature	73
10.2.6. Timestamp	75
10.2.7. Nonce	75
10.2.8. Trust Anchor	76
10.2.9. Certificate	76
10.2.10. Hashed Message Authentication Code (HMAC)	76
10.2.11. Neighbor Synchronization	78

10.2.12. Interface Attributes	79
10.2.13. Traffic Selector	84
10.2.14. Geo Coordinates	86
10.2.15. Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Message	87
10.2.16. PIM-SM Message	87
10.2.17. Fragmentation Report (FRAGREP)	88
10.2.18. ICMPv6 Error	89
10.2.19. Proxy/Server Departure	90
11. Address Mapping - Multicast	91
12. Multilink Conceptual Sending Algorithm	91
12.1. Multiple OMNI Interfaces	92
12.2. Client-Proxy/Server Loop Prevention	93
13. Router Discovery and Address/Prefix Delegation	93
13.1. Client-Proxy/Server Window Synchronization	104
13.2. Router Discovery in IP Multihop and IPv4-Only Networks	106
13.3. DHCPv6-based Prefix Registration	110
14. Secure Redirection	111
15. Proxy/Server Resilience	111
16. Detecting and Responding to Proxy/Server Failures	112
17. Transition Considerations	112
18. OMNI Interfaces on Open Internetworks	113
19. Time-Varying MNPs	115
20. Error Messages	115
21. IANA Considerations	115
21.1. Protocol Numbers	116
21.2. IEEE 802 Numbers	116
21.3. IPv4 Special-Purpose Address	116
21.4. IANA OUI Ethernet Numbers	116
21.5. Overlay Multilink Network (OMNI) Interface Registry Group	117
21.5.1. OMNI Option Sub-Types (New Registry)	117
21.5.2. OMNI Node Identification ID-Types (New Registry)	117
21.5.3. OMNI Geo Coordinates Types (New Registry)	118
21.6. Additional Considerations	118
22. Security Considerations	119
23. Implementation Status	120
24. Document Updates	120
25. Acknowledgements	121
26. References	122
26.1. Normative References	122
26.2. Informative References	125
Appendix A. IPv6 Compatible Addresses	135
Appendix B. IPv6 ND Message Authentication and Integrity	136
Appendix C. VDL Mode 2 Considerations	137
Appendix D. Client-Proxy/Server Isolation Through Link-Layer Address Mapping	138

Appendix E. IPv6 ND Virtual Interface Model	138
Appendix F. IPv4 as an OAL Encapsulation Service	139
Appendix G. Change Log	139
Author's Address	141

1. Introduction

Air/land/sea/space mobile nodes (e.g., aircraft of various configurations, terrestrial vehicles, seagoing vessels, space systems, enterprise wireless devices, pedestrians with cellphones, etc.) configure mobile routers with multiple interface connections to wireless and/or wired-line data links. These data links often have diverse performance, cost and availability properties that can change dynamically according to mobility patterns, flight phases, proximity to infrastructure, etc. The mobile router acts as a Client of a network-based Mobility Service (MS) by configuring a virtual interface over its underlay interface data link connections.

Each Client configures a virtual network interface (termed the "Overlay Multilink Network (OMNI) Interface") as a thin layer over its underlay interfaces which may themselves connect to virtual or physical links. The OMNI interface is therefore the only interface abstraction exposed to the IP layer and behaves according to the Non-Broadcast, Multiple Access (NBMA) interface principle, while each underlay interface appears as a link layer communication channel in the architecture. The OMNI interface appears as a "virtual Ethernet (veth)" interface to the IP layer and internally employs the "OMNI Adaptation Layer (OAL)" to ensure that original IP packets are adapted to diverse underlay interfaces with heterogeneous properties.

The OMNI interface connects to a virtual overlay known as the "OMNI link". The OMNI link spans one or more Internetworks that may include private-use infrastructures (e.g., enterprise networks, operator networks, etc.) and/or the global public Internet itself. Together, OMNI and the OAL provide the foundational elements required to support the "6 M's of Modern Internetworking", including:

1. Multilink - a Client's ability to coordinate multiple diverse underlay interfaces as a single logical unit (i.e., the OMNI interface) to achieve the required communications performance and reliability objectives.
2. Multinet - the ability to span the OMNI link over a segment routing topology with multiple diverse administrative domain network segments while maintaining seamless end-to-end communications between mobile Clients and correspondents such as air traffic controllers, fleet administrators, etc.

3. Mobility - a Client's ability to change network points of attachment (e.g., moving between wireless base stations) which may result in an underlay interface address change, but without disruptions to ongoing communication sessions with peers over the OMNI link.
4. Multicast - the ability to send a single network transmission that reaches multiple Clients belonging to the same interest group, but without disturbing other Clients not subscribed to the interest group.
5. Multihop - a mobile Client peer-to-peer relaying capability useful when multiple forwarding hops between peers may be necessary to reach a target peer or an infrastructure access point connection to the OMNI link.
6. (Performance) Maximization - the ability to exchange large packets between peers without loss due to a link size restriction, and to adaptively adjust packet sizes to maintain the best performance profile for each independent traffic flow.

Client OMNI interfaces coordinate with the MS and/or OMNI peer nodes through secure IPv6 Neighbor Discovery (ND) control message exchanges [RFC4861]. The MS consists of a distributed set of service nodes (including Proxy/Servers and other infrastructure elements) that also configure OMNI interfaces. Automatic Extended Route Optimization (AERO) in particular provides a companion MS compatible with the OMNI architecture [I-D.templin-6man-aero3]. AERO discusses details of ND message based multilink forwarding, route optimization, mobility management, and multinet traversal while the fundamental aspects of OMNI link operation are discussed in this document.

Each OMNI interface provides a multilink nexus for exchanging inbound and outbound traffic flows via selected underlay interfaces. The IP layer sees the OMNI interface as a point of connection to the OMNI link. Each OMNI link assigns one or more associated Mobility Service Prefixes (MSPs), which are typically IP Global Unicast Address (GUA) prefixes. The MS then delegates Mobile Network Prefixes (MNPs) taken from an MSP to Client end systems as PI address blocks. Clients in local domains also obtain PA addresses from internal/external Stable Network Prefixes (SNPs) assigned to Proxy/Servers that connect the local domain to the global topology per [RFC6296]. If there are multiple OMNI links, the IP layer will see multiple OMNI interfaces.

Clients receive SNP addresses and optionally also MNP prefix delegations through IPv6 ND control message exchanges with Proxy/Servers over MANETs, Access Networks (ANETs) and/or open Internetworks (INETs). Clients sub-delegate MNPs to downstream-

attached End-user Networks (ENETs) independently of the underlay interfaces selected for upstream data transport. Each Client acts as a fixed or mobile router on behalf of ENET peers, and uses OMNI interface control messaging to coordinate with Proxy/Servers and/or other Clients. The Client iterates its control messaging over each of the OMNI interface's (M)ANET/INET underlay interfaces in order to register each interface with the MS (see Section 13). The Client can also provide (proxied) multihop forwarding services for a recursively extended chain of other Clients and end systems connected via downstream-attached *NETs.

Each Proxy/Server on the link delegates SNP GUA addresses from its unique IPv6 prefix to Clients via the Dynamic Host Configuration Protocol for IPv6 (DHCPv6) service. DHCPv6 messaging is carried as IPv6 ND message extensions since each Proxy/Server provides the combined services of a DHCPv6 Server and IPv6 router. Since the DHCPv6 service running on each Proxy/Server provides assurance against address duplication within its own prefix, Clients need not test the IPv6 SNP GUA addresses they receive for uniqueness. Clients instead regard each Proxy/Server as the address delegation authority and designated router for a distinct OMNI link segment.

Clients may connect to multiple distinct OMNI links within the same OMNI domain by configuring multiple OMNI interfaces, e.g., omni0, omni1, omni2, etc. Each OMNI interface is configured over a distinct set of underlay interfaces and provides a nexus for Safety-Based Multilink (SBM) operation. The IP layer applies SBM routing to select a specific OMNI interface, then the selected OMNI interface applies Performance-Based Multilink (PBM) internally to select appropriate underlay interfaces. Applications select SBM topologies based on IP layer Segment Routing [RFC8402], while each OMNI interface orchestrates PBM internally based on OAL Multinet traversal.

OMNI provides a link model suitable for a wide range of use cases. For example, the International Civil Aviation Organization (ICAO) Working Group-I Mobility Subgroup is developing a future Aeronautical Telecommunications Network with Internet Protocol Services (ATN/IPS) and has issued a liaison statement requesting IETF adoption [ATN] in support of ICAO Document 9896 [ATN-IPS]. The IETF IP Wireless Access in Vehicular Environments (ipwave) working group has further included problem statement and use case analysis for OMNI in [RFC9365]. Still other communities of interest include AEEC, RTCA Special Committee 228 (SC-228) and NASA programs that examine commercial aviation, Urban Air Mobility (UAM) and Unmanned Air Systems (UAS). Pedestrians with handheld mobile devices using 5G/6G wireless services, home and small office networks, enterprise networks and many others represent additional large classes of potential OMNI users.

This document specifies the transmission of original IP packets and control messages over OMNI interfaces. The operation of both IP protocol versions (i.e., IPv4 [RFC0791] and IPv6 [RFC8200]) is specified as the network layer data plane, while OMNI interfaces use IPv6 ND messaging in the control plane independently of the data plane protocol(s). OMNI interfaces also provide an adaptation layer based on encapsulation and fragmentation over heterogeneous underlay interfaces as an OAL sublayer between L3 and L2. OMNI and the OAL are specified in detail throughout the remainder of this document.

2. Terminology

The terminology in the normative references applies; especially, the terms "link" and "interface" are the same as defined in the IPv6 [RFC8200] and IPv6 Neighbor Discovery (ND) [RFC4861] specifications. This document assumes the following IPv6 ND control plane message types: Router Solicitation (RS), Router Advertisement (RA), Neighbor Solicitation (NS), Neighbor Advertisement (NA), unsolicited NA (uNA) and Redirect. AERO further introduces new IPv6 ND pseudo-message types Multilink Initiate (MI), Multilink Respond (MR) and Multilink Control (MC) with formats identical to the standard uNA message but with different Code values. These messages are used to control adaptation layer functions only and are not exposed to the network layer. See [I-D.templin-6man-aero3] for a full specification of MI/MR/MC.

The terms "All-Routers multicast", "All-Nodes multicast" and "Subnet-Router anycast" are the same as defined in [RFC4291]. Also, IPv6 ND state names, variables and constants including REACHABLE, ReachableTime and REACHABLE_TIME are the same as defined in [RFC4861].

The term "IP" is used to refer collectively to either Internet Protocol version (i.e., IPv4 [RFC0791] or IPv6 [RFC8200]) when a specification at the layer in question applies equally to either version.

The terms (Proxy/)Client and Proxy/Server are intentionally capitalized to denote an instance of a particular node type that also configures an OMNI interface and engages the OMNI Adaptation Layer.

The terms "application layer (L5 and higher)", "transport layer (L4)", "network layer (L3)", "(data) link layer (L2)" and "physical layer (L1)" are used consistently with common Internetworking terminology, with the understanding that reliable delivery protocol users of UDP are considered as transport layer elements. The OMNI specification further defines an "adaptation layer" positioned below the network layer but above the link layer, which may include

physical links and Internet- or higher-layer tunnels. A (network) interface is a node's attachment to a link (via L2), and an OMNI interface is therefore a node's attachment to an OMNI link (via the adaptation layer).

The terms "IP jumbogram", "advanced jumbo (AJ)" and "IP parcel" refer to alternative packet formats discussed in [I-D.templin-6man-parcels2] and [I-D.templin-intarea-parcels2].

The following terms are defined within the scope of this document:

GUA, ULA, LLA, MLA

A Globally-Unique (GUA), Unique-Local (ULA) or Link-Local (LLA) Address per the IPv6 addressing architecture [RFC4193] [RFC4291], or a Multilink-Local Address (MLA) per [I-D.templin-6man-mla]. IPv4 prefixes other than those reserved for special purposes [RFC6890] are also considered as GUA prefixes.

L3

The Network layer in the OSI network model. Also known as "layer 3", "IP layer", etc.

L2

The Data Link layer in the OSI network model. Also known as "layer 2", "link layer", "sub-IP layer", etc.

Adaptation layer

An encapsulation mid-layer that adapts L3 to a diverse collection of L2 underlay interfaces and their encapsulations. (No layer number is assigned, since numbering was an artifact of the legacy reference model that need not carry forward in the modern architecture.) The adaptation layer sees the network layer as "L3" and sees all link layer encapsulations as "L2 encapsulations", which may include UDP, IP and true link layer (e.g., Ethernet, etc.) headers.

Access Network (ANET)

a connected network region (e.g., an aviation radio access network, corporate enterprise network, satellite service provider network, cellular operator network, residential WiFi network, etc.) that connects Clients to the rest of the OMNI link. Physical and/or data link level security is assumed (sometimes referred to as "protected spectrum" for wireless domains). ANETs such as private enterprise networks and ground domain aviation service networks often provide multiple secured IP hops between the Client's physical point of connection and the nearest Proxy/Server.

Mobile Ad-hoc NETwork (MANET)

a connected ANET region for which links often have undetermined connectivity properties, lower layer security services cannot always be assumed and multihop forwarding between Clients acting as MANET routers may be necessary. Clients nested deeply within a MANET often require adaptation layer proxy services from "upstream" peer Proxy/Clients located progressively nearer the MANET border.

Internetwork (INET)

a connected network region with a coherent IP addressing plan that provides transit forwarding services between ANETs and/or OMNI nodes that coordinate with the Mobility Service over unprotected media. Since physical and/or data link level security cannot always be assumed, security must be applied by the network and/or higher layers if necessary. The global public Internet itself is an example.

End-user Network (ENET)

a simple or complex "downstream" network tethered to a Client as a single logical unit that travels together. The ENET could be as simple as a single link connecting a single host, or as complex as a large network with many links, routers, bridges and end user devices. The ENET provides an "upstream" link for arbitrarily many low-, medium- or high-end devices dependent on the Client for their upstream connectivity, i.e., as Internet of Things (IoT) entities. The ENET can also support a recursively-descending chain of additional Clients such that the ENET of an upstream Client is seen as the ANET of a downstream Client.

***NET**

a "wildcard" term used when a given specification applies equally to all MANET/ANET/INET cases. From the Client's perspective, *NET interfaces are "upstream" interfaces that connect the Client to the Mobility Service, while ENET interfaces are "downstream" interfaces that the Client uses to connect downstream ENETs and/or other Clients. Local communications between correspondents within the same *NET can often be conducted based on IPv6 ULAs [RFC4193] or MLAs [I-D.templin-6man-mla].

underlay interface

a *NET or ENET interface over which an OMNI interface is configured. The OMNI interface is seen as an L3 interface by the network layer, and each underlay interface is seen as an L2 interface by the OMNI interface. The underlay interface either connects directly to the physical communications media or coordinates with another node where the physical media is hosted.

MANET Interface

a node's underlay interface to a local network with indeterminant neighborhood properties over which multihop relaying may be necessary. All MANET interfaces used by AERO/OMNI are IPv6 interfaces and therefore must configure a Maximum Transmission Unit (MTU) no smaller than the IPv6 minimum MTU (1280 octets) even if lower-layer fragmentation is needed.

OMNI link

a Non-Broadcast, Multiple Access (NBMA) virtual overlay configured over one or more INETs and their connected (M)ANETs/ENETs. An OMNI link may comprise multiple distinct "segments" joined by "bridges" the same as for any link; the addressing plans in each segment may be mutually exclusive and managed by different administrative entities. Proxy/Servers and other infrastructure elements extend the link to support communications between Clients as single-hop neighbors.

OMNI link segment

a Proxy/Server and all of its constituent Clients within any attached *NETs is considered as a leaf OMNI link segment, with each leaf interconnected via links and "bridge" nodes in intermediate OMNI link segments. When the *NETs of multiple leaf segments overlap (e.g., due to network mobility), they can combine to form larger *NETs with no changes to Client-to-Proxy/Server relationships. The OMNI link consists of the concatenation of all OMNI link leaf and intermediate segments as a loop-free spanning tree.

OMNI interface

a node's virtual Ethernet (veth) interface to an OMNI link, and configured over one or more underlay interfaces. If there are multiple OMNI links in an OMNI domain, a separate OMNI interface is configured for each link. The OMNI interface configures a Maximum Transmission Unit (MTU) and an Effective MTU to Receive (EMTU_R) the same as any interface. The OMNI interface assigns an LLA the same as for any IPv6 interface and assigns an MLA for adaptation layer addressing over its underlay networks. The OMNI interface further assigns any unicast or anycast ULA/GUA addresses acquired through address autoconfiguration. Since OMNI interface addresses are managed for uniqueness, OMNI interfaces do not require Duplicate Address Detection (DAD) and therefore set the administrative variable 'DupAddrDetectTransmits' to zero [RFC4862].

OMNI Adaptation Layer (OAL)

an OMNI interface sublayer service that encapsulates original IP packets admitted into the interface in an IPv6 header and/or

subjects them to fragmentation and reassembly. The OAL is also responsible for generating MTU-related control messages as necessary, and for providing addressing context for OMNI link SRT traversal. The OAL presents a new layer in the Internet architecture known simply as the "adaptation layer". The OMNI link is an example of a limited domain [RFC8799] at the adaptation layer although its segments may be joined over open Internetworks at L2.

OMNI Option

a pseudo IPv6 ND option providing multilink parameters for the OMNI interface as specified in Section 10. The OMNI option is appended to the end of an IPv6 ND message during OAL encapsulation such that it appears immediately following the final message option.

(OMNI) Client

a network platform/device mobile router or end system that configures one or more OMNI interfaces over distinct sets of underlay interfaces grouped as logical OMNI link units. The Client coordinates with the Mobility Service via upstream networks over *NET interfaces, and may also serve as a Proxy for other Clients on downstream *NETs. The Client's upstream network interface addresses and performance characteristics may change over time (e.g., due to node mobility, link quality, etc.) while other Clients on downstream networks can engage the (upstream) Client as a Proxy.

(OMNI) Proxy/Server

a segment routing topology edge node that configures an OMNI interface and connects Clients to the Mobility Service. As a server, the Proxy/Server responds directly to some Client IPv6 ND messages. As a proxy, the Proxy/Server forwards other Client IPv6 ND messages to other Proxy/Servers and Clients. As a router, the Proxy/Server provides a forwarding service for ordinary data messages that may be essential in some environments and a last resort in others. Proxy/Servers at (M)ANET boundaries configure both an (M)ANET downstream interface and *NET upstream interface, while INET-based Proxy/Servers configure only an INET interface. All Proxy/Servers configure a Stable Network Prefix (SNP) and manage 1x1 mappings of internal ULAs and external GUAs according to [RFC6296].

First-Hop Segment (FHS) Proxy/Server

a Proxy/Server connected to the source Client's *NET that forwards OAL packets sent by the source into the segment routing topology. FHS Proxy/Servers allocate Provider-Aggregated (Proxy/Server-Aggregated) addresses to Clients within their local networks. FHS

Proxy/Servers also act as intermediate forwarding systems to facilitate RS/RA-based Provider-Independent Prefix Delegation exchanges between Clients and Mobility Anchor Point (MAP) Proxy/Servers.

Last-Hop Segment (LHS) Proxy/Server

a Proxy/Server connected to the target Client's *NET that forwards OAL packets received from the segment routing topology to the target.

Mobility Anchor Point (MAP) Proxy/Server

a Proxy/Server selected by the Client that provides a designated router service for any *NET underlay networks that register the Client's Mobile Network Prefix (MNP). Since all Proxy/Servers provide equivalent services, Clients normally select the first FHS Proxy/Server they coordinate with to serve as the MAP. However, the MAP can instead be any available Proxy/Server for the OMNI link, i.e., and not necessarily one of the Client's FHS Proxy/Servers. This flexible arrangement supports a fully distributed mobility management service.

Segment Routing Topology (SRT)

a multinet forwarding region configured over one or more INETs between the FHS Proxy/Server and LHS Proxy/Server. The SRT spans the OMNI link on behalf of communicating peer nodes using segment routing in a manner outside the scope of this document (see: [I-D.templin-6man-aero3]).

Mobility Service (MS)

a mobile routing service that tracks Client movements and ensures that Clients remain continuously reachable even across mobility events. The MS consists of the set of all Proxy/Servers plus all other OMNI link supporting infrastructure nodes. Specific MS details are out of scope for this document, with an example found in [I-D.templin-6man-aero3].

Mobility Service Prefix (MSP)

an aggregated IP GUA prefix (e.g., 2001:db8::/32, 2002:192.0.2.0::/40, etc.) assigned to the OMNI link and from which more-specific Mobile and Stable Network Prefixes (MNPs/SNPs) are delegated, where IPv4 MSPs are represented as "6to4 prefixes" per [RFC3056]. OMNI link administrators typically obtain MSPs from an Internet address registry, however private-use prefixes can also be used subject to certain limitations (see: Section 8). OMNI links that connect to the global Internet advertise their MSPs to their interdomain routing peers.

Mobile Network Prefix (MNP)

a longer IP prefix delegated from an MSP (e.g., 2001:db8:1000:2000::/56, 2002:192.0.2.8::/46, etc.) and assigned to a Client. Clients receive MNPs from MAP Proxy/Servers and sub-delegate them to routers in downstream ENETs.

Stable Network Prefix (SNP)

a ULA/GUA IP prefix pair assigned to one or more Proxy/Servers that connect local *NET Client groups to the rest of the OMNI link. Clients request address delegations from the SNP that can be used to support global and local-scoped communications. Clients communicate internally within *NET groups using IPv6 ULAs assigned in 1:1 correspondence to SNP GUAs made visible to external peers through IP network address/prefix translation [RFC6145][RFC6146][RFC6147][RFC6296].

Foreign Network Prefix (FNP)

a global IP prefix not covered by a MSP and assigned to a link or network outside of the OMNI domain.

Subnet Router Anycast (SRA) Address

An IPv6 address taken from an FNP/MNP/SNP in which the remainder of the address beyond the prefix is set to the value "all-zeros". For example, the SRA for 2001:db8:1::/48 is simply 2001:db8:1:: (i.e., with the 80 least significant bits set to 0). For IPv4, the IPv6 SRA corresponding to the IPv4 prefix 192.0.2.0/24 is 2002:192.0.2.0::/40 per [RFC3056].

Provider-Aggregated (PA) Address

a ULA/GUA address pair delegated to a Client from an FHS Proxy/Server SNP is considered Provider-Aggregated (PA) or "Proxy/Server-Aggregated". The Client either assigns the GUA PA address to its own OMNI interface or allows the FHS Proxy/Server to supply the address via Network Prefix Translation for IPv6 (NPTv6) [RFC6296].

Provider-Independent (PI) Address

a GUA allocated from an MNP delegated to a Client via a MAP Proxy/Server is considered Provider-Independent (PI) or "Proxy/Server-Independent". The Client assigns PI addresses to (downstream) ENET interfaces and can also sub-delegate the MNP to downstream ENET nodes.

original IP packet

a whole IP packet or fragment admitted into the OMNI interface by the network layer prior to OAL encapsulation/fragmentation, or an IP packet delivered to the network layer by the OMNI interface following OAL reassembly/decapsulation.

OAL packet

an original IP packet encapsulated in an OAL IPv6 header with an IPv6 Extended Fragment Header extension that includes an 8-octet (64-bit) OAL Identification value. Each OAL packet is then subject to fragmentation by the source and reassembly by the destination.

OAL fragment

a portion of an OAL packet following fragmentation but prior to L2 encapsulation, or following L2 decapsulation but prior to OAL reassembly.

(OAL) atomic fragment

an OAL packet that can be forwarded without fragmentation, but still includes an IPv6 Extended Fragment Header with an 8-octet (64-bit) OAL Identification value and with Index and More Fragments both set to 0.

(L2) carrier packet

an encapsulated OAL fragment following L2 encapsulation or prior to L2 decapsulation. OAL sources and destinations exchange carrier packets over underlay interfaces, and may be separated by one or more OAL intermediate systems. OAL intermediate systems may perform re-encapsulation on carrier packets by removing the L2 headers of the first hop network and replacing them with new L2 headers for the next hop network. Carrier packets may themselves be subject to fragmentation and reassembly in L2 underlay networks at a layer below the OAL. Carrier packets sent over unsecured paths use OMNI protocol L2 encapsulations, while those sent over secured paths use L2 security encapsulations such as IPsec [RFC4301]. (The term "carrier" honors agents of the service postulated by [RFC1149] and [RFC6214].)

OAL source

an OMNI interface acts as an OAL source when it encapsulates original IP packets to form OAL packets, then performs OAL fragmentation and encapsulation to create carrier packets which may themselves be subject to fragmentation at lower layers. Every OAL source is also an OMNI link ingress.

OAL destination

an OMNI interface acts as an OAL destination when it decapsulates carrier packets (while reassembling first, if necessary), then performs OAL reassembly/decapsulation to derive the original IP packet. Every OAL destination is also an OMNI link egress.

OAL intermediate system

an OMNI interface acts as an OAL intermediate system when it decapsulates carrier packets received from a first segment to obtain the original OAL packet/fragment, then re-encapsulates in new L2 headers and sends these new carrier packets into the next segment. OAL intermediate systems decrement the Hop Limit in OAL packets/fragments during forwarding, and discard the OAL packet/fragment if the Hop Limit reaches 0. OAL intermediate systems do not decrement the TTL/Hop Limit of the original IP packet, which can only be updated by the network and higher layers. OAL intermediate systems along the path explicitly addressed by the OAL IPv6 Destination (e.g., Proxys, etc.) are regarded as "endpoint" intermediate systems while those not explicitly addressed (e.g., MANET routers, AERO Gateways, etc.) are regarded as "transit" intermediate systems.

Multilink

a Client OMNI interface's manner of managing multiple diverse *NET underlay interfaces as a single logical unit. The OMNI interface provides a single unified interface to the network layer, while underlay interface selections are performed on a per-flow basis considering traffic selectors such as Traffic Class, Flow Label, application policy, signal quality, cost, etc. Multilink selections are coordinated in both the outbound and inbound directions based on source/target underlay interface pairs.

Multinet

an intermediate system's manner of spanning multiple diverse IP Internetwork and/or private enterprise network "segments" through OAL encapsulation. Multiple diverse Internetworks (such as the global public IPv4 and IPv6 Internets) can serve as transit segments in an end-to-end OAL forwarding path through intermediate system concatenation of SRT network segments. This OAL concatenation capability provides benefits such as supporting IPv4/IPv6 transition and coexistence, joining multiple diverse operator networks into a cooperative single service network, etc. See: [I-D.templin-6man-aero3] for further information.

Multihop

an iterative relaying of carrier packets between Clients over an OMNI underlay interface technology (such as omnidirectional wireless) without support of fixed infrastructure. Multihop services entail Client-to-Client relaying within a Mobile/Vehicular Ad-hoc Network (MANET/VANET) for Vehicle-to-Vehicle (V2V) communications and/or for Vehicle-to-Infrastructure (V2I) "range extension" where Clients within range of communications infrastructure elements provide forwarding services for other Clients.

Mobility

any action that results in a change to a Client underlay interface address. The change could be due to, e.g., a handover to a new wireless base station, loss of link due to signal fading, an actual physical node movement, etc.

Safety-Based Multilink (SBM)

A means for ensuring fault tolerance through redundancy by connecting multiple OMNI interfaces within the same domain to independent routing topologies (i.e., multiple independent OMNI links).

Performance Based Multilink (PBM)

A means for selecting one or more underlay interface(s) for carrier packet transmission and reception within a single OMNI interface.

OMNI Domain

The set of all SBM/PBM OMNI links that collectively provides services for a common set of MSPs. All OMNI links within the same domain configure, advertise and respond to the SRA address(es) corresponding to the MSP(s) assigned to the domain.

flow

a sequence of packets sent from a particular source to a particular unicast, anycast, or multicast destination that a node desires to label as a flow. The 3-tuple of the Flow Label, Source Address and Destination Address fields enable efficient IPv6 flow classification. The IPv6 Flow Label Specification is observed per [RFC6437] [RFC6438].

AERO Flow Information Base (AFIB)

A multilink forwarding table on each OAL source, destination and intermediate system that includes AERO Flow Vectors (AFV) with both next hop forwarding instructions and context for reconstructing compressed headers for specific underlay interface pairs used to transport flows from a source to a destination. See: [I-D.templin-6man-aero3] for further discussion.

AERO Flow Vector (AFV)

An AFIB entry that includes soft state for each underlay interface pairwise communication flow from source to destination. AFVs are identified by an AFV Index (AFVI) paired with the previous hop L2 address, with the pair established based on an IPv6 ND solicitation and solicited IPv6 ND advertisement response. The AFV also caches underlay interface pairwise Identification sequence number parameters to support carrier packet filtering. See: [I-D.templin-6man-aero3] for further discussion.

AERO Flow Vector Index (AFVI)

A 2-octet or 4-octet integer value supplied by a previous hop OAL node when it requests a next hop OAL node to create an AFV. (The AFVI is always processed as a 4-octet value, but compressed headers may omit the 2 most significant octets when they encode the value 0.) The next hop OAL node caches the AFVI and L2 address supplied by the previous hop as header compression/decompression state for future OAL packets with compressed headers. The previous hop OAL node must ensure that the AFVI values it assigns to the next hop via a specific underlay interface are distinct and reused only after their useful lifetimes expire. The special value 0 means that no AFVI is asserted.

(OMNI) L2 encapsulation

the OMNI protocol encapsulation of OAL packets/fragments in an outer header or headers to form carrier packets that can be routed within the scope of the local *NET underlay network partition. The OAL node that performs encapsulation is known as the "L2 source" while the OAL node that performs decapsulation is known as the "L2 destination"; both OAL end and intermediate systems can also act as an L2 source or destination. Common L2 encapsulation combinations include UDP, IP and/or Ethernet using a port/protocol/type number for OMNI.

L2 address (L2ADDR)

an address that appears in the OMNI protocol L2 encapsulation for an underlay interface and also in IPv6 ND message OMNI options. L2ADDR can be either an IP address for IP encapsulations or an IEEE EUI address [EUI] for direct data link encapsulation. (When UDP/IP encapsulation is used, the UDP port number is considered an ancillary extension of the IP L2ADDR.)

OAL Fragment Size (OFS)

the current OAL source fragmentation size for a given flow which must be no smaller than 1024 octets and should be no larger than 65279 octets to allow sufficient space for OAL and L2 encapsulations. (OFS pertains to the fragment payload immediately following the fragment header; if OAL extension headers are included following the first fragment header a slightly larger minimum OFS may be necessary to accommodate maximum-sized packets.) Each OAL source maintains OFS in an AERO Flow Vector (AFV) for each independent flow. The OAL source discovers larger OFS sizes through dynamic probing the same as defined for Maximum Packet Size (MPS) probing per Section 4.4 of [RFC8899] and should adaptively maintain the best possible OFS for each flow according to current network conditions.

3. Requirements

OMNI interfaces maintain the same Conceptual Data Structures as for any IPv6 interface, including the Neighbor Cache, Destination Cache, Prefix List and Default Router List [RFC4861]. The same as for any IPv6 interface, different routers on the link may advertise different prefixes. The OMNI interface must therefore ensure that any addresses configured from the prefixes and assigned to the interface are associated with the correct default routers.

OMNI interfaces should limit the size of their IPv6 ND control plane messages (plus any original IP packet attachments) to the adaptation layer path MTU which may be as small as the minimum IPv6 link MTU minus encapsulation overhead. If there are sufficient OMNI parameters and/or IP packet attachments that would exceed this size, the OMNI interface should forward the information as multiple smaller IPv6 ND messages and the recipient accepts the union of all information received. This allows the messages to travel without loss due to a size restriction over secured control plane paths that include IPsec tunnels [RFC4301], secured direct point-to-point links and/or unsecured paths that require an authentication signature.

Client and Proxy/Server OMNI interfaces maintain per-destination state in Destination Cache Entries (DCEs) as a level of indirection into per-neighbor state in Neighbor Cache Entries (NCEs). The function of these caches and the IPv6 ND Protocol Constants defined in Section 10 of [RFC4861] apply for this document.

The L3, adaptation and (virtual) L2 layers each include distinct packet Identification numbering spaces. The adaptation layer employs an 8-octet Identification numbering space that is distinct from L3/L2 spaces, with an Identification value appearing in an IPv6 Extended Fragment Header [I-D.templin-6man-ipid-ext2] or an OMNI Compressed Header (OCH) (see: Section 6.5) in each adaptation layer encapsulation.

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.

4. Overlay Multilink Network (OMNI) Interface Model

The IP layer sees the OMNI interface as a virtual Ethernet (veth) interface configured over one or more underlay interfaces, which may be physical (e.g., an aeronautical radio link, a cellular wireless link, etc.) or virtual (e.g., an internet-layer or higher-layer "tunnel"). The OMNI interface architectural layering model is the same as in [RFC5558][RFC7847], and augmented as shown in Figure 1. The network layer therefore sees the OMNI interface as a single L3 interface nexus for multiple underlay interfaces that appear as L2 communication channels in the architecture.

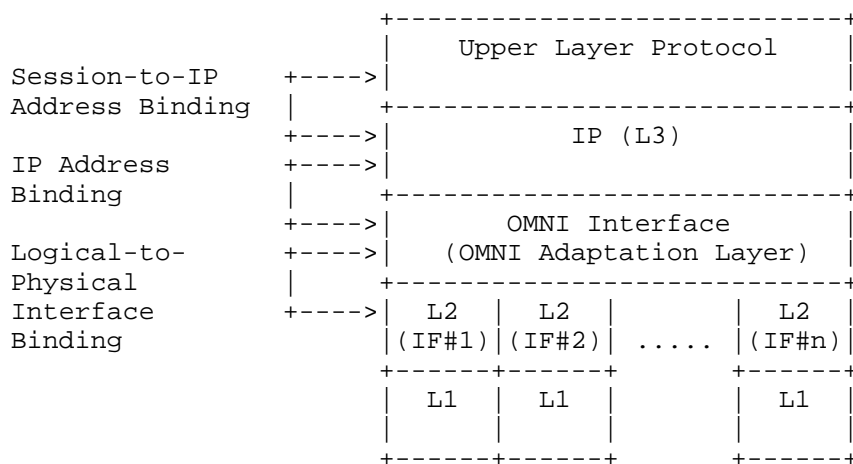


Figure 1: OMNI Interface Architectural Layering Model

Each underlay interface provides an L2/L1 abstraction according to one of the following models:

- * (M)ANET interfaces connect to a (M)ANET that is separated from the open INET by Proxy/Servers. The (M)ANET interface may be either on the same link segment as a Proxy/Server, or separated from a Proxy/Server by multiple adaptation layer and/or L2 hops. (Note that NATs may appear internally within a (M)ANET or on the Proxy/Server itself and may require NAT traversal the same as for the INET case.)
- * INET interfaces connect to an INET either natively or through IP Network Address Translators (NATs). Native INET interfaces have global IP addresses that are reachable from any INET correspondent. NATed INET interfaces typically configure private IP addresses and connect to a private network behind one or more NATs with the outermost NAT providing INET access.

- * ENET interfaces connect a Client's downstream-attached networks, where the Client provides forwarding services for ENET end system communications to remote peers. An ENET may be as simple as a small IoT sub-network that travels with a mobile Client to as complex as a large private enterprise network that the Client connects to a larger *NET. Downstream-attached Clients engage the ENET as a *NET and engage the (upstream) Client as a Proxy.
- * VPN interfaces use security encapsulations (e.g. IPsec tunnels) over underlay networks to connect Client, Proxy/Server or other critical infrastructure nodes. VPN interfaces provide security services at lower layers of the architecture (L2/L1), with securing properties similar to Direct point-to-point interfaces.
- * Direct point-to-point interfaces securely connect Clients, Proxy/Servers and/or other critical infrastructure nodes over physical or virtual media that does not transit any open Internetwork paths. Examples include a line-of-sight link between a remote pilot and an unmanned aircraft, a fiberoptic link between gateways, etc.

The OMNI interface forwards original IP packets from the network layer using the OMNI Adaptation Layer (OAL) (see: Section 5) as an encapsulation and fragmentation sublayer service. This "OAL source" then further encapsulates the resulting OAL packets/fragments in underlay network headers (e.g., UDP/IP, IP-only, Ethernet-only, etc.) to create L2 encapsulated "carrier packets" for fragmentation and transmission over underlay interfaces. The target OMNI interface then receives the carrier packets from underlay interfaces and performs L2 decapsulation.

If the resulting OAL packets/fragments are addressed to itself, the OMNI interface performs reassembly/decapsulation as an "OAL destination" and delivers the original IP packet to the network layer. If the OAL packets/fragments are addressed to another node, the OMNI interface instead re-encapsulates them in new underlay network L2 headers as an "OAL intermediate system" then forwards the resulting carrier packets over an underlay interface. The OAL source and OAL destination are seen as "neighbors" on the OMNI link, while OAL intermediate systems provide a virtual bridging service that joins the segments of a (multinet) Segment Routing Topology (SRT).

The OMNI interface transports carrier packets over either secured or unsecured underlay interfaces to access the secured/unsecured OMNI link spanning trees as discussed further throughout the document. Carrier packets that carry control plane messages over secured underlay interfaces use secured L2/L1 services such as IPsec, direct encapsulation over secured point-to-point links, etc. Carrier

packets that carry data plane messages over unsecured underlay interfaces instead use L2 encapsulations appropriate for public or private Internetworks and are subject for the following sections.

The OMNI interface and its OAL can forward original IP packets over underlay interfaces while including/omitting various lower layer encapsulations including OAL, UDP, IP and (ETH)ernet or other link layer header. The network layer can also engage underlay interfaces directly while bypassing the OMNI interface entirely when necessary. This architectural flexibility may be beneficial for underlay interfaces (e.g., some aviation data links) for which encapsulation overhead is a primary consideration. OMNI interfaces that send original IP packets directly over underlay interfaces without invoking the OAL can only reach peers located on the same OMNI link segment. Source Clients can instead use the OAL to coordinate with target Clients in the same or different OMNI link segments by sending initial carrier packets to a First-Hop Segment (FHS) Proxy/Server. The FHS Proxy/Server then sends the carrier packets into the SRT spanning tree, which transports them to a Last-Hop Segment (LHS) Proxy/Server for the target Client.

The OMNI interface encapsulation/decapsulation layering possibilities are shown in Figure 2 below. An imaginary vertical lines drawn between the Network Layer at the top of the figure and Underlay Interfaces at the bottom of the figure then allowed to slide horizontally either to the right or left illustrates the various encapsulation/decapsulation layering combination possibilities. Common combinations include IP-only (i.e., direct access to underlay interfaces with or without using the OMNI interface), IP/IP, IP/UDP/IP, IP/UDP/IP/ETH, IP/OAL/UDP/IP, IP/OAL/UDP/ETH, etc.

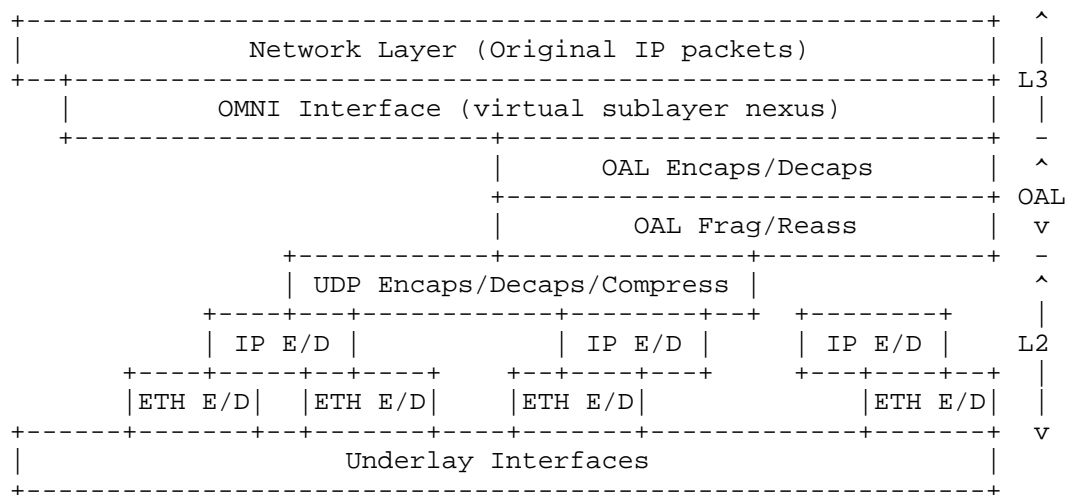


Figure 2: OMNI Interface Layering

The OMNI/OAL model gives rise to a number of opportunities:

- * Clients coordinate with the MS and receive both SNP addresses and MNP delegations through IPv6 ND control plane message exchanges with Proxy/Servers. Since GUA and ULA addresses are managed for uniqueness, no Duplicate Address Detection (DAD) or Multicast Listener Discovery (MLD) messaging is necessary over the OMNI interface.
- * underlay interfaces on the same L2 link segment as a Proxy/Server do not require any L3 addresses (i.e., not even link-local) in environments where communications are coordinated entirely over the OMNI interface.
- * as underlay interface properties change (e.g., link quality, cost, availability, etc.), any active interface can be used to update the profiles of multiple additional interfaces in a single message. This allows for timely adaptation and service continuity under dynamically changing conditions.
- * coordinating underlay interfaces in this way allows them to be represented in a unified MS profile with provisions to support the "6 M's of Modern Internetworking".
- * header compression and path MTU determination is conducted on a per-flow basis, with each flow adapting to the best performance profiles and path selections.
- * exposing a single virtual interface abstraction to the network layer allows for multilink operation (including QoS based link selection, carrier packet replication, load balancing, etc.) at L2 while still permitting L3 traffic shaping based on, e.g., Traffic Class, Flow Label, etc.
- * the OMNI interface supports multinet traversal over the SRT when communications across different administrative domain network segments are necessary. This mode of operation would not be possible via direct communications over the underlay interfaces themselves.
- * the OAL supports lossless and adaptive path MTU mitigations not available for communications directly over the underlay interfaces themselves. The OAL supports "packing" of multiple original IP payload packets within a single OAL "composite packet" and also supports transmission of IP packets of all sizes up to and including (advanced) jumbograms.

- * the OAL assigns per-packet Identification values that allow for adaptation/link layer reliability and data origin authentication.
- * L3 sees the OMNI interface as a point of connection to the OMNI link; if there are multiple OMNI links, L3 will see multiple OMNI interfaces.
- * Multiple independent OMNI interfaces can be used for increased fault tolerance through Safety-Based Multilink (SBM), with Performance-Based Multilink (PBM) applied within each interface.
- * Multiple independent OMNI links can be joined together into a single link without requiring renumbering of infrastructure elements, since the GUAs/ULAs assigned by Proxy/Servers of the different links will be mutually exclusive.
- * OMNI provides robust support for both Provider-Aggregated (PA) and Provider-Independent (PI) addressing resulting in a versatile service for all Client use cases.
- * The concept of OMNI endpoint intermediate systems allows for logical partitioning within MANETs without requiring address aggregation. Instead, MANET routing within each partition is based on MLA "host routes" that are not redistributed into other partitions. Each partition connects via multiple interface Proxy/Clients in a hierarchy of partitions on the path to an FHS Proxy/Server.

Figure 3 depicts the architectural model for a source Client with an attached ENET connecting to the OMNI link via multiple independent *NETs. The Client's OMNI interface forwards adaptation layer IPv6 ND solicitation messages over available *NET underlay interfaces using any necessary L2 encapsulations. The IPv6 ND messages traverse the *NETs until they reach an FHS Proxy/Server (FHS#1, FHS#2, ..., FHS#n), which returns an IPv6 ND advertisement message and/or forwards a proxied version of the message over the SRT to an LHS Proxy/Server near the target Client (LHS#1, LHS#2, ..., LHS#m). The Hop Limit in IPv6 ND messages is not decremented due to encapsulation; hence, the source and target Client OMNI interfaces appear to be attached to a shared NBMA link.

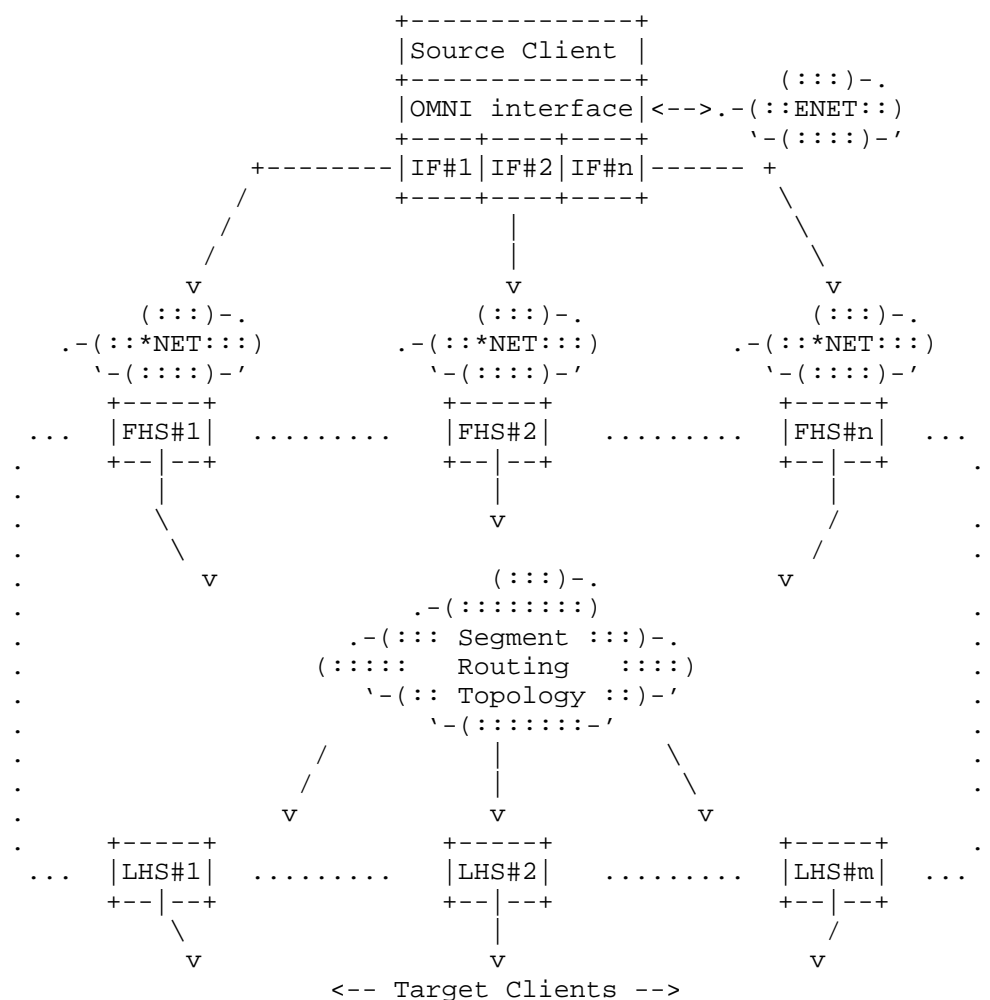


Figure 3: Source/Target Client Coordination over the OMNI Link

After the initial IPv6 ND message exchange, the source Client (as well as any nodes on its attached ENETs) can send carrier packets to the target Client via the OMNI interface. OMNI interface multilink services will forward the carrier packets via FHS Proxy/Servers for the correct underlay *NETs. The FHS Proxy/Server then re-encapsulates the carrier packets and forwards them over the SRT which delivers them to an LHS Proxy/Server, and the LHS Proxy/Server in turn re-encapsulates and forwards them to the target Client. (Note that when the source and target Client are on the same SRT segment, the FHS and LHS Proxy/Servers may be one and the same.)

Mobile Clients select a MAP Proxy/Server (not shown in the figure), which will often be one of their FHS Proxy/Servers but may also be any Proxy/Server on the OMNI link. Clients then register all of their *NET underlay interfaces with the MAP Proxy/Server via per interface FHS Proxy/Servers in a pure proxy role. The MAP Proxy/Server then provides a designated router that advertises the Client's MNPs into the OMNI link routing system, and the Client can quickly migrate to a new MAP Proxy/Server if the former becomes unresponsive.

Clients therefore use Proxy/Servers as bridges into the SRT to reach OMNI link correspondents via a spanning tree established in a manner outside the scope of this document. Proxy/Servers forward critical MS control messages via the secured spanning tree and forward other messages via the unsecured spanning tree (see Security Considerations). When AERO route optimization is applied, Clients can instead forward directly to correspondents in the same SRT segment to reduce Proxy/Server and/or Gateway load.

Note: Original IP packets sent into an OMNI interface will receive consistent consideration according to their size as discussed in the following sections, while those sent directly over underlay interfaces that exceed the underlay network path MTU are dropped with an ordinary ICMP Packet Too Big (PTB) message returned. These PTB messages are subject to loss the same as for any non-OMNI IP interface [RFC2923].

5. OMNI Interface Maximum Transmission Unit (MTU)

The OMNI interface observes the link nature of tunnels, including the Maximum Transmission Unit (MTU), Effective MTU to Send (EMTU_S), Effective MTU to Receive (EMTU_R) and the role of fragmentation and reassembly [I-D.ietf-intarea-tunnels]. The OMNI interface is configured over one or more underlay interfaces as discussed in Section 4, where underlay links and network paths may have diverse MTUs. OMNI interface considerations for accommodating original IP packets of various sizes are discussed in the following sections.

IPv6 underlay interfaces are REQUIRED to configure a minimum MTU of 1280 octets and a minimum EMTU_R of 1500 octets [RFC8200]. Therefore, the minimum IPv6 path MTU is 1280 octets since routers on the path are not permitted to perform network fragmentation even though the destination is required to reassemble more. The network therefore MUST forward original IP packets as large as 1280 octets without generating an IPv6 Path MTU Discovery (PMTUD) Packet Too Big (PTB) message [RFC8201].

IPv4 underlay interfaces are REQUIRED to configure a minimum MTU of 68 octets [RFC0791] and a minimum EMTU_R of 576 octets [RFC0791][RFC1122]. However, links that configure small MTUs are likely to have low-end performance and occur only at the extreme network edges while higher-performance interior network links should configure MTUs no smaller than 1280 octets and EMTU_Rs no smaller than 1500 octets [RFC3819].

The OMNI interface itself sets an "unlimited" MTU of $(2^{32} - 1)$ octets. The network layer therefore unconditionally admits all original IP packets into the OMNI interface, where the adaptation layer accommodates them if possible according to their size. The OAL source then invokes adaptation layer encapsulation/fragmentation services to transform all original IP packets no larger than 65535 octets into OAL packets/fragments. The OAL source then applies L2 encapsulation to form carrier packets and finally forwards the carrier packets via underlay interfaces.

When the OAL source performs IPv6 encapsulation and fragmentation (see: Section 6), the Payload Length field limits the maximum-sized original IP packet that the OAL can accommodate while applying IPv6 fragmentation to $(2^{16} - 1) = 65535$ octets (i.e., not including the OAL encapsulation header lengths). The OAL source is also permitted to forward packets larger than this size as a best-effort delivery service if the L2 path can accommodate them through "jumbo-in-jumbo" encapsulation (see: [I-D.templin-6man-parcels2]); otherwise, the OAL source discards the packet and arranges to return a PTB "hard error" to the original source (see: Section 6.9).

Each OMNI interface therefore sets a minimum EMTU_R of 65535 octets (plus the length of the OAL encapsulation headers), and each OAL destination must consistently either accept or reject still larger whole packets that arrive over any of its underlay interfaces according to their size. If an underlay interface presents a whole packet larger than the OAL destination is prepared to accept (e.g., due to a buffer size restriction), the OAL destination discards the packet and arranges to return a PTB "hard error" to the OAL source which in turn forwards a translated PTB to the original source (see: Section 6.9).

6. The OMNI Adaptation Layer (OAL)

The OMNI interface forwards original IP packets from the network layer for transmission over one or more underlay interfaces. The OMNI Adaptation Layer (OAL) acting as the OAL source then replaces the virtual Ethernet header with an IPv6 encapsulation header to form OAL packets. OAL source fragmentation then breaks the packets into IPv6 fragments suitable for L2 encapsulation and transmission as

carrier packets.

The carrier packets then traverse one or more underlay networks spanned by OAL intermediate systems in the SRT. Each successive OAL intermediate system then re-encapsulates by removing the L2 headers of the first underlay network and appending L2 headers appropriate for the next underlay network. (This process supports the multinet concatenation capability needed for joining multiple diverse networks.) Following any forwarding by OAL intermediate systems, the carrier packets eventually arrive at the OAL destination.

When the OAL destination receives the carrier packets, it discards the L2 headers and reassembles the resulting OAL fragments into an OAL packet as described in Section 6.3. The OAL destination next replaces the OAL packet IPv6 encapsulation header with a virtual Ethernet header to obtain the original IP packet for delivery to the network layer via the OMNI interface. The OAL source may be either the source Client or its FHS Proxy/Server, while the OAL destination may be either the LHS Proxy/Server or the target Client. Proxy/Servers (and SRT Gateways per [I-D.templin-6man-aero3]) may also serve as OAL intermediate systems.

The OAL presents an OMNI sublayer abstraction similar to ATM Adaptation Layer 5 (AAL5). Unlike AAL5 which performs segmentation and reassembly with fixed-length 53-octet cells over ATM networks, however, the OAL uses IPv6 encapsulation, fragmentation and reassembly with larger variable-length cells over heterogeneous networks. Detailed operations of the OAL are specified in the following sections.

6.1. OAL Source Encapsulation and Fragmentation

When the network layer forwards an original IP packet into the OMNI interface, it either sets the TTL/Hop Limit for locally-generated packets or decrements the TTL/Hop Limit according to standard IP forwarding rules. The OAL source next creates an "OAL packet" by replacing the virtual Ethernet header with an IPv6 encapsulation header per [RFC2473]. The OAL source sets the IPv6 encapsulation header Version to "OMNI-IP6" (see: Section 6.2) and Next Header to TBD1 (see: IANA Considerations).

When the OAL source performs IPv6 encapsulation, it sets the IPv6 header "Flow Label" as specified in [RFC6438]. The OAL source next copies the "Type of Service/Traffic Class Differentiated Service Code Point (DSCP)" [RFC2474][RFC2983] and "Explicit Congestion Notification (ECN)" [RFC3168] values in the original packet's IP header into the corresponding fields of the OAL IPv6 header.

For original IP packets with DSCP '111111' (including ordinary network control/data plane messages), the OAL source resets the OAL IPv6 encapsulation header DSCP to '110111'. The OAL source instead sets the IPv6 encapsulation header DSCP to '111111' for adaptation layer control plane messages that must be processed by all OAL intermediate systems on the path including both endpoints and transits. These DSCP values belong to the "Pool 2 Experimental and Local Use (EXP/LU)" range reserved in [RFC2474] and correspond to Network/Internetwork Control precedence in [RFC0791].

The OAL source next sets the IPv6 header Payload Length to the length of the original IP packet and sets Hop Limit to a value that is sufficiently large to support loop-free forwarding over multiple concatenated OAL intermediate hops. The OAL source next selects OAL IPv6 Source and Destination Addresses associated with its own OMNI interface and the OMNI interface of the target. (These are MLA addresses that correspond to the virtual Ethernet source and destination MAC addresses as maintained in a per neighbor address mapping cache for the interface - see: Section 8.)

The OAL source next inserts any necessary extension headers following the IPv6 header as specified in Section 6.4. For OAL data plane packets, the source first inserts any per-fragment extension headers (e.g., Hop-by-Hop, Routing, etc.) then inserts an IPv6 Extended Fragment Header (see: [I-D.templin-6man-ipid-ext2]) with an 8-octet (64-bit) OAL packet Identification. Note that the extension header insertions could cause the IPv6 Payload Length to exceed 65535 octets by a small amount when the original IP packet is (nearly) the maximum length.

The OAL source then source-fragments the OAL packet if necessary according to an OAL Fragment Size (OFS) maintained in AERO Flow Vectors (AVFs) for each independent flow. (The OAL source encapsulates payloads that are no larger than the OFS and original IP packets larger than 65535 octets as "atomic fragments".) OAL fragments prepared by the source must not be fragmented further by OAL intermediate systems on the path to the OAL destination.

OAL packets that contain original IP packets no larger than 65535 octets are subject to OAL source fragmentation using the IPv6 Extended Fragment Header (EFH) fragmentation specification [I-D.templin-6man-ipid-ext2] with the exception that the IPv6 Payload Length may exceed 65535 by at most the length of the extension headers. For each independent flow, the OAL source MUST set OFS to a size no smaller than 1024 octets and thereafter adjust OFS according to dynamic network control message feedback. The OAL source SHOULD limit OFS to a size no larger than 65279 octets unless it has assurance that the path can accommodate a larger size. (Note: the

minimum size ensures that OAL fragments can be accommodated over any potential combination of IPv4/IPv6 underlay network paths where transit for smaller sizes is assured without probing, while the maximum size observes the 65535 octet limitation for conventional IP packets.)

When the OAL source performs fragmentation, it SHOULD produce the minimum number of fragments under current OFS constraints. The fragments produced MUST be non-overlapping and the portion of each non-final fragment following the IPv6 Extended Fragment Header MUST be equal in length while that of the final fragment MAY be smaller and MUST NOT be larger.

For each consecutive OAL fragment beginning with the first, the OAL source then writes a monotonically-increasing "ordinal" value between 0 and 63 in the Extended Fragment Header Index field. Specifically, the OAL source writes the ordinal value '0' for the first fragment, '1' for the first non-first fragment, '2' for the next, '3' for the next, etc. up to the final fragment. The final fragment may assign an ordinal as large as '63' such that at most 64 fragments are possible. During a network path change, an OAL intermediate system may apply further OAL fragmentation to produce minimum-length (sub-)fragments. The OAL destination will then reassemble these (sub-)fragments then combine each reassembled fragment with all other fragments of the same OAL packet and return rate-limited indications to inform the OAL source that the path has changed.

The OAL source finally encapsulates the fragments in L2 headers to form carrier packets for transmission over underlay interfaces, while retaining the fragments and their ordinal numbers (i.e., #0, #1, #2, etc.) for a brief period to support adaptation layer retransmissions (see: Section 6.8). OAL fragment and carrier packet formats are shown in Figure 4.






Figure 4: OAL Fragments and Carrier Packets

After establishing AFV state in the forward path for a given flow, the OAL source dynamically manages the per-flow OFS by continually probing the forward path to the OAL destination beginning with a size no smaller than 1024 octets and increasing to progressively larger sizes per [RFC8899]. In this process, the OAL source acts as a datagram packetization layer for the flow when it applies OAL encapsulation, fragmentation and header compression.

The OAL source creates a probe by setting the P flag in the Type 1 OMNI Compressed Header (OCH1) (see: Section 6.5) of a probe packet for the flow. For probes that advance the OFS to a larger size, the probe packet can include discard data (e.g., an IP packet with Next Header/Protocol set to 59 ("No Next Header"), a UDP packet with service port number set to 9 ("discard"), etc.) or live protocol data with null padding. For probes that confirm the current OFS, the probe packet can instead entirely include live protocol data. The OAL source then admits the probe for L2 encapsulation and transmission.

When the OAL destination receives the probe, it returns an OAL-encapsulated secured control message to the OAL source with an OMNI option that includes an ICMPv6 Error sub-option. The OAL destination sets the ICMPv6 message Type to 2 ("Packet Too Big") and Code to "MTU Probe Reply" (see: [I-D.templin-6man-ipid-ext2]), then sets MTU to the size of the probe message minus the difference in size between the OAL/IP full headers and the OCH1 header. The OAL destination then copies the leading 512 octets of the probe beginning with the full OAL and IP headers (i.e., replacing the OCH1 header) into the ICMPv6 message body. The OAL destination then returns the secured control message to the OAL source without marking it for examination by OAL intermediate systems.

When the OAL source receives the secured control message, it can determine that the message did not originate from an attacker. The OAL source can then tentatively advance OFS for this flow to the larger size reported in the ICMPv6 MTU option (minus OAL encapsulation overhead) but should maintain an ongoing stream of additional probes for the flow to confirm the current OFS and/or to advance to still larger OFS values. The OAL source may additionally receive MTU soft error feedback from an OMNI destination or intermediate system and should compensate accordingly as discussed in Section 6.9.

6.2. OAL L2 Encapsulation and Re-Encapsulation

The OAL source or intermediate system next encapsulates each OAL fragment (with either full or compressed headers) in L2 encapsulation headers to create a carrier packet. The OAL source or intermediate system (i.e., the L2 source) includes a UDP header as the innermost sublayer if NATs and/or filtering middleboxes might occur on the path. Otherwise, the L2 source includes a full/compressed IP header and/or an actual link layer header (e.g., such as for Ethernet-compatible links) as the innermost sublayer. The L2 source also appends any additional encapsulation sublayer headers necessary (e.g., IPsec AH/ESP, jumbo-in-jumbo encapsulation, etc.).

The L2 source encapsulates the OAL information immediately following the innermost L2 sublayer header. The L2 source next interprets the first 4 bits following the L2 headers as a Type field that determines the type of OAL header that follows. The OAL source sets Type to (OMNI-IP6) for an uncompressed IPv6 OMNI Full Header or (OMNI-OCH1/2) for an OMNI Compressed Header (OCH1/2) as specified in Section 6.5. Other Type values may also appear as specified in Section 6.5.

The OAL node prepares the L2 encapsulation headers for OAL packets/fragments as follows:

- * For UDP/IP encapsulation, the L2 source sets the UDP source port to 8060 (i.e., the port number reserved for AERO/OMNI). When the L2 destination is a Proxy/Server or Gateway, the L2 source sets the UDP Destination Port to 8060; otherwise, the L2 source sets the UDP Destination Port to its cached port number value for the peer. The L2 source next sets the UDP Length the same as specified in [I-D.ietf-tsvwg-udp-options].
- * The L2 source then sets the IP {Protocol, Next Header} to '17' (the UDP protocol number) and sets the {Total, Payload} Length the same as specified in the base IP protocol specifications for ordinary IP packets (see: [RFC0791], [RFC8200] and [I-D.ietf-tsvwg-udp-options]). The L2 source then continues to set the remaining IP header fields as discussed below.
- * For raw IP encapsulation, the L2 source sets the IP {Protocol, Next Header} to TBD1 (see: IANA Considerations) and sets the {Total, Payload} Length the same as specified in [RFC0791] or [RFC8200]. The L2 source then continues to set the remaining IP header fields as discussed below.
- * For IPsec AH/ESP encapsulation, the L2 source sets the appropriate IP or UDP header to indicate AH/ESP then sets the AH/ESP Next Header field to TBD1 the same as for raw IP encapsulation.
- * For direct encapsulations over Ethernet-compatible links, the L2 source prepares an Ethernet Header with EtherType set to TBD2 (see: Section 21.2) (see: Section 7).
- * For OAL packet/fragment encapsulations over secured underlay interface connections to the secured spanning tree, the L2 source applies any L2 security encapsulations according to the protocol (e.g., IPsec). These secured carrier packets are then subject to lower layer security services.

When an L2 source includes a UDP header, it SHOULD calculate and include a UDP checksum in carrier packets with full OAL headers to prevent mis-delivery and/or detect IPv4 reassembly corruption; the L2 source MAY set UDP checksum to 0 (disabled) in carrier packets with compressed OAL headers (see: Section 6.5) or when reassembly corruption is not a concern. If the L2 source discovers that a path is dropping carrier packets with UDP checksums disabled, it should supply UDP checksums in future carrier packets sent to the same L2 destination. If the L2 source discovers that a path is dropping carrier packets that do not include a UDP header, it should include a UDP header in future carrier packets.

When an L2 source sends carrier packets with compressed OAL headers and with UDP checksums disabled, mis-delivery due to corruption of the AFVI is possible but unlikely since the corrupted index would somehow have to match valid state in the (sparsely-populated) AERO Flow Information Base (AFIB). In the unlikely event that a match occurs, an OAL destination may receive carrier packets that contain a mis-delivered OAL fragment but can immediately reject any with incorrect Identifications. If the Identification value is somehow accepted, the OAL destination may submit the mis-delivered OAL fragment to the reassembly cache where it will most likely be rejected due to incorrect reassembly parameters. If a reassembly that includes the mis-delivered OAL fragment somehow succeeds (or, for atomic fragments) the OAL destination will verify any included checksums to detect corruption. Finally, any spurious data that somehow eludes all prior checks will be detected and rejected by end-to-end upper layer integrity checks. See: [RFC6935] [RFC6936] for further discussion.

For UDP/IP or IP-only L2 encapsulations, when the L2 source is also the OAL source it next copies the DSCP, ECN and Flow Label values from the OAL header into the L2 header. The L2 source then sets the L2 IP TTL/Hop Limit the same as for any host (i.e., it does not copy the Hop Limit value from the OAL header) and finally sets the IP Source and Destination Addresses to direct the carrier packet to the next OAL hop. For carrier packets subject to re-encapsulation, the OAL intermediate system removes the L2 header(s) then prepares to act as the L2 source for the next hop.

The L2 source first decrements the OAL header Hop Limit and discards the OAL packet/fragment if the value reaches 0. Otherwise, the L2 source copies the DSCP value from OAL IPv6 header into the next segment L2 encapsulation header while setting the next segment L2 IP Source and Destination Addresses the same as above. The L2 source then copies the ECN value from the previous segment L2 encapsulation header into both the OAL full/compressed header and the next segment L2 encapsulation header.

The L2 source then prepares to forward the carrier packets to the next OAL intermediate system or destination. For L2 encapsulations over IPv4, if the carrier packet is no larger than 1280 octets the L2 source sets the IPv4 Don't Fragment (DF) bit to 0 and includes a suitable IPv4 Identification value; otherwise, the OAL source sets DF to 1. This ensures that all IPv4 carrier packets no larger than 1280 octets will be delivered to the L2 destination even if a small amount of fragmentation occurs in the path (see: [RFC3819] for IPv4 link MTU expectations according to their performance characteristics).

For IPv4 carrier packets that set DF to 1 and for all IPv6 carrier packets, delivery is best-effort according to the available path MTU in the spirit of [RFC2473] and [RFC4213]. Since carrier packet transmissions are not within the scope of an explicit tunnel required to pass the IPv6 minimum MTU, however, there is no need for the L2 source to apply L2 source fragmentation since the 1024 octet minimum OFS is operationally assured over all IPv4 and IPv6 paths. The L2 source should therefore ignore any ICMPv6 Packet Too Big or IPv4 Fragmentation Needed messages returned from the network in response to any of its large carrier packet transmissions since the OAL source engages in active probing per [RFC8899].

The L2 source then sends the resulting carrier packets over one or more underlay interfaces. Underlay interfaces often connect directly to physical media on the local platform (e.g., an aircraft with a radio frequency link, a laptop computer with WiFi, etc.), but in some configurations the physical media may be hosted on a separate Local Area Network (LAN) node. In that case, the OMNI interface can establish a Layer-2 VLAN or a point-to-point tunnel (at a layer below the underlay interface) to the node hosting the physical media. The OMNI interface may also apply encapsulation at the underlay interface layer (e.g., as for a tunnel virtual interface) such that carrier packets would appear "double-encapsulated" on the LAN; the node hosting the physical media in turn removes the LAN encapsulation prior to transmission or inserts it following reception. Finally, the underlay interface must monitor the node hosting the physical media (e.g., through periodic keepalives) so that it can convey up-to-date Interface Attribute information to the OMNI interface.

6.3. Reassembly and Decapsulation

For both IPv4 and IPv6, OAL intermediate systems and destinations MUST configure an L2 minimum EMTU_R of 1500 octets on all unsecured underlay interfaces. (Secured underlay interfaces instead use an EMTU_R specific to the L2 security service such as IPsec.) OAL intermediate systems and destinations are permitted to configure a larger L2 EMTU_R in order to pass larger unfragmented carrier packets, but need not reassemble more than 1500.

OAL destinations MUST configure an adaptation layer EMTU_R of 65535 octets to support reassembly of fragmented OAL packets of all sizes. OAL nodes must further recognize and honor the extended Identifications included in the IPv6 Extended Fragment Header [I-D.templin-6man-ipid-ext2].

When an OMNI interface processes a carrier packet received on an underlay interface, it copies the ECN value from the L2 encapsulation headers into the OAL header but does not copy the DSCP value from the

L2 encapsulation headers into the OAL header according to the differentiated services pipe model for tunnels [RFC2983]. The OMNI interface next discards the L2 encapsulation headers and examines the OAL header of the enclosed OAL packet/fragment according to the value in the Type field as discussed in Section 6.2

If the OAL packet/fragment is addressed to a different node, the OMNI interface (acting as an OAL intermediate system) decrements the OAL Hop Limit as discussed in Section 6.2 then performs L2 encapsulation and forwards the resulting carrier packet. If the OAL packet/fragment is addressed to itself, the OMNI interface (acting as an OAL destination) accepts or drops based on the (Source, Destination, Identification)-tuple.

The OAL destination next drops all ordinal OAL non-first fragments that would overlap or leave "holes" with respect to other ordinal fragments already received. The OAL destination updates a checklist of accepted ordinal fragments of the same OAL packet but admits all accepted fragments into the reassembly cache.

During reassembly at the OAL destination, the reassembled OAL packet may exceed 65535 by a small amount equal to the size of the OAL encapsulation extension headers. The OAL destination does not write this (too-large) value into the OAL header Payload Length field, but instead remembers the value during reassembly. When reassembly is complete, the OAL destination finally replaces the OAL IPv6 encapsulation header with a virtual Ethernet header. The OAL destination's OMNI interface then delivers the original IP packet to the network layer. The original IP packet may therefore be as large as 65535 octets.

When an OAL path traverses an IPv6 network with routers that perform adaptation layer forwarding based on full IPv6 headers with OAL addresses, the OAL intermediate system at the head of the IPv6 path forwards the OAL packet/fragment the same as an ordinary IPv6 packet without decapsulating and delivering to the network layer. Once within the IPv6 network, these OAL packets/fragments may traverse arbitrarily-many IPv6 hops before arriving at an OAL intermediate system which may again encapsulate the OAL packets/fragments as carrier packets for transmission over underlay interfaces.

Note: carrier packets often traverse paths with underlying links that use integrity checks such as CRC-32 which provide adequate hop-by-hop integrity assurance for payloads up to ~9K octets [CRC]. However, other paths may traverse links (such as fragmenting tunnels over IPv4 - see: [RFC4963]) that do not include adequate checks.

6.4. OMNI-Encoded IPv6 Extension Headers

The IPv6 specification [RFC8200] defines extension headers that follow the base IPv6 header, while Upper Layer Protocols (ULPs) are specified in other documents. Each extension header present is identified by a "Next Header" octet in the previous (extension) header and encodes a "Next Header" field in the first octet that identifies the next extension header or ULP instance. The OMNI specification supports encoding of IPv6 extension header chains immediately following the OMNI L2 UDP, IP or Ethernet header even if the L2 IP protocol version is IPv4. In all cases, the length of the IPv6 extension header chain is limited by [I-D.ietf-6man-eh-limits].

The OAL source prepares an OMNI extension header chain by setting the first 4 bits of the first IPv6 extension header in the chain to a Type value for the extension header itself immediately following the OMNI L2 protocol header. The source then sets the next 4 bits to a Next value that identifies either a terminating ULP or the next extension header in the chain. The source then sets the first 8 bits of each subsequent IPv6 extension header in the chain to the standard Next Header encoding as shown in Figure 5:




Figure 5: OMNI Extension Header Chains

The following Type/Next values are currently defined:

- 0 (OMNI-RES1) - Reserved for experimentation.
- 1 (OMNI-OCH1) - OMNI Compressed Header, Type 1 per Section 6.5.

- 2 (OMNI-OCH2) - OMNI Compressed Header, Type 2 per Section 6.5.
- 3 (OMNI-RES2) - Reserved for experimentation.
- 4 (OMNI-IP4) - IPv4 header per [RFC0791].
- 5 (OMNI-HBH) - Hop-by-Hop Options per Section 4.3 of [RFC8200].
- 6 (OMNI-IP6) - IPv6 header per [RFC8200].
- 7 (OMNI-RH) - Routing Header per Section 4.4 of [RFC8200].
- 8 (OMNI-FH) - Fragment Header per Section 4.5 of [RFC8200].
- 9 (OMNI-DO) - Destination Options per Section 4.6 of [RFC8200].
- 10 (OMNI-AH) - Authentication Header per [RFC4302].
- 11 (OMNI-ESP) - Encapsulating Security Payload per [RFC4303].
- 12 (OMNI-NNH) - No Next Header per Section 4.7 of [RFC8200].
- 13 (OMNI-TCP) - TCP Header per [RFC9293].
- 14 (OMNI-UDP) - UDP Header per [RFC0768].
- 15 (OMNI-ULP) - Upper Layer Protocol shim (see below).

Entries OMNI-OCH1 through OMNI-AH in the above list follow the convention that the OMNI Type/Version appears in the first 4 bits of the extension header (or IP header) itself. Conversely, entries OMNI-ESP through OMNI-UDP represent commonly-used ULPs which do not encode a Type/Version in the first 4 bits.

Entries OMNI-HBH, OMNI-RH, OMNI-FH, OMNI-DO and OMNI-AH represent true IPv6 extension headers encoded for OMNI, which may be chained. Source and destination processing of OMNI extension headers follows exactly per their definitions in the normative references, with the exception of the special (Type, Next) coding in the first 8 bits of the first extension header.

When a ULP not found in the above table immediately follows the OMNI L2 UDP, IP or Ethernet header, the source includes a 2-octet "Type 1 ULP Shim" before the ULP where both the first 4 bit (Type) and next 4 bit (Next) fields encode the special value 15 (OMNI-ULP). The source then includes a Next Header field that encodes the IP protocol number of the ULP. The source then includes the ULP data immediately after the shim as shown in Figure 6.

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|Type=15|Next=15|   Next Header   |   Upper Layer Protocol   |~
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 6: OMNI Upper Layer Protocol (ULP) Shim (Type 1)

When a ULP "OMNI-(N)" found in the above table immediately follows the OMNI L2 UDP, IP or Ethernet header, the source includes a 1-octet "Type 2 ULP Shim" before the ULP where the first 4 bits encode the special Type value 15 (OMNI-ULP) and the next 4 bits encode the Next ULP type "N" taken from the table above. The source then includes the ULP data immediately after the shim as shown in Figure 7.

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|Type=15| Next=N|               Upper Layer Protocol               |~
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 7: OMNI Upper Layer Protocol (ULP) Shim (Type 2)

When a ULP not found in the above table follows a first OMNI extension header, the source sets the extension header Next field to OMNI-ULP (15) and includes a 1-octet "Type 3 ULP Shim" that encodes the IP protocol number for the Next Header of the ULP data that follows as shown in Figure 8.

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|   Next Header   |               Upper Layer Protocol               |~
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 8: OMNI Upper Layer Protocol (ULP) Shim (Type 3)

When a ULP "OMNI-(N)" found in the above table follows a first OMNI extension header, the source sets the extension header Next field to the ULP Type "N" and does not include a shim. The ULP then begins immediately after the first OMNI extension header.

When a ULP of any kind follows a non-first OMNI extension header, the source sets the extension header Next Header field to the IP protocol number for the ULP and does not include a shim. The ULP then begins immediately after the non-first OMNI extension header.

Note: The L2 UDP header (when present) is logically considered as the first L2 extension header in the chain. If an Advanced Jumbo extension header is also present, its Jumbo Payload length includes the length of the L2 UDP header.

Note: After a node parses the extension header chain, it changes the "Type/Next" field in the first extension header back to the correct "Next Header" value before processing the first extension header.

6.5. OMNI Full and Compressed Headers

OAL sources that send OAL packets with full OMNI IPv6 Headers include a Segment Routing Header (SRH) as an extension per [RFC8754]. The Segment List elements include the MLAs of the OAL end systems themselves, the MLAs of any edge network partition border OAL intermediate systems and the SNP SRA GUAs of OAL intermediate systems in the global Internetwork. Client end systems discover the Segment List elements in their RS/RA exchanges with Internetworking Proxy/Servers, where each partition border OAL intermediate system in the RS message forward path records its MLA before forwarding to the next partition border OAL intermediate system.

The SRH is followed by an IPv6 Extended Fragment Header to support segment-by-segment forwarding based on an AERO Flow Information Base (AFIB) in each OAL intermediate system. OAL sources, intermediate systems and destinations establish header compression state in the AFIB through IPv6 ND control message exchanges. After an initial control message exchange, OAL nodes can apply OMNI Header Compression to significantly reduce header overhead.

OAL nodes apply header compression in order to avoid transmission of redundant data found in the original IP packet and OAL encapsulation headers; the resulting compressed headers are often significantly smaller than the original IP packet header itself even when OAL encapsulation is applied. Header compression is limited to the OAL IPv6 encapsulation header plus extensions along with the base original IP packet header; it does not extend to include any extension headers of the original IP packet which appear as upper layer payload immediately following the compressed headers.

Each OAL node establishes AFIB soft state entries known as AERO Flow Vectors (AFVs) which support both OAL packet/fragment forwarding and OAL/IPv6 header compression/decompression. The FHS OAL sources references each AFV by an AERO Flow Vector Index (AFVI) which in conjunction with the previous hop L2ADDR provides compression/decompression and next hop forwarding context.

When an OAL node sends carrier packets that contain OAL packets/fragments to a next hop, it includes a full IPv6 header with an SRH containing segment addressing information followed by an Extended Fragment Header. The first 4 bits following the L2 headers must encode the Type OMNI-IP6 to signify that an uncompressed IPv6 header (plus any extensions) is present.

When AFV state is available, the OAL source should omit significant portions of the OAL header (plus extensions) and original IP packet header by applying OMNI header compression. For OAL first fragments (including atomic fragments), the OAL source uses OMNI Compressed Header, Type 1 (OCH1) Format (a) as shown in Figure 9:

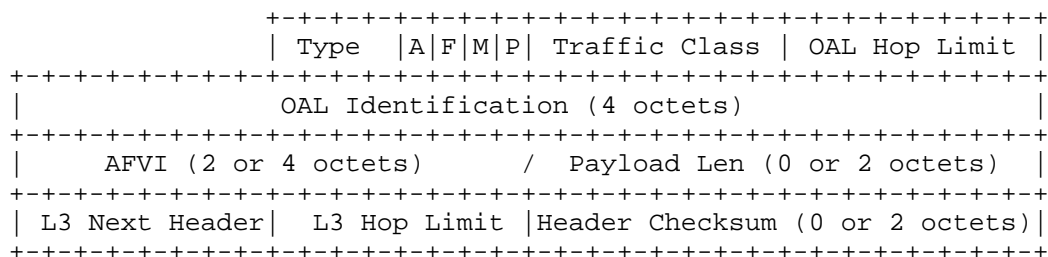


Figure 9: OMNI Compressed Header (OCH1) Format (a)

The format begins with a 4-bit Type followed by 4 flag bits followed by an 8-bit Traffic Class (copied into the OAL header from the original IP packet header) followed by an 8-bit (OAL) Hop Limit.

The header next includes the 4 least significant octets of the OAL Identification followed by a 2/4-octet AFVI according to whether the (A) flag is set to 0/1, respectively. The format then includes a 2-octet Payload Length only if the L2 header does not include a length field. The format finally includes the Next Header and Hop Limit values from the original (L3) IP packet header, plus a 2-octet Header Checksum only for IPv4 original packets. (Note that these values represent compression of the original IP packet header plus the OAL IPv6 header along with its SRH and Extended Fragment Header in a unified concatenation.)

The OAL node sets Type to OMNI-OCH1, sets Hop Limit to the uncompressed OAL header Hop Limit and sets the ECN bits in the Traffic Class field the same as for an uncompressed IP header. The OAL node next sets (F)ormat to 0 then sets (M)ore Fragments the same as for an uncompressed Extended Fragment Header. The OAL node finally sets the L3 Next Header and Hop Limit fields to the values that would appear in the uncompressed original IP header; the OAL node also includes a 2-octet Header Checksum for IPv4 original packets, or omits the Header Checksum for IPv6 original packets.

The payload of the OAL first fragment (i.e., beginning after the original IP header) is then included immediately following the OCH1 header, and the L2 header length field (if present) is reduced by the difference in length between the compressed and full-length headers. If the L2 header includes a length field, the OAL destination can

determine the payload length by examining the L2 header; otherwise, the OCH1 header itself includes a 2-octet Payload Length field that encodes the length of the packet payload that follows the OCH1. Note that OAL first fragments (and atomic packets) are logically considered ordinal fragment 0.

When the OAL source needs to probe the OAL Fragment Size (OFS) for a given flow, it sets the (P)robe flag and includes a probe message of the desired size following the OCH1 header. Upon receipt, the OAL destination returns a secured control message reply to the OAL source. When the OAL source receives the control message, it can either maintain its current OFS for this flow or advanced to a larger OFS according to the probe size.

For OAL non-first fragments (i.e., those with non-zero Index), the OAL uses OMNI Compressed Header, Type 1 (OCH1) Format (b) as shown in Figure 10:

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| Type  | A|F|M|R|Res|   Index   | Traffic Class | OAL Hop Limit |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Identification (4 octets)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      AFVI (2 or 4 octets)      /   Payload Len (0 or 2 octets)   |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 10: OMNI Compressed Header (OCH1) Format (b)

The format begins with a 4-bit Type followed by 4 flags followed by a 2-bit Reserved field (set to 0) followed by a 6-bit ordinal fragment Index. All other fields up to and including the Payload Length (if present) are included the same as for an OCH1 first fragment.

The OAL node sets Type to OMNI-OCH1, sets Hop Limit to the uncompressed OAL header Hop Limit value, sets (Index, (A)FVI, (M)ore Fragments, Identification) to their appropriate values as a non-first fragment and sets (F)ormat to 1. In the process, the OAL Node sets Index to a monotonically increasing ordinal value beginning with 1 for the first non-first fragment, 2 for the second non-first fragment, 3 for the third non-first fragment, etc., up to at most 63 for the final fragment.

The OAL non-first fragment body is then included immediately following the OCH1 header, and the L2 header length field (if present) is reduced by the difference in length between the compressed headers and full-length original IP header with OAL IPv6 header plus extensions. The OAL destination will then be able to determine the Payload Length by examining the L2 header length field if present; otherwise by examining the 2-octet OCH1 Payload Length the same as for first fragments.

The OCH1 Format (a) is used for all original IPv6 packets that do not include a Fragment Header as well as for original IPv4 packets that set IHL to 5, DF to 1 and (MF; Fragment Offset) to 0 (the OCH1 Format (b) is used for non-first fragments in all IP protocol cases).

For other "non-atomic" original IP packets and first fragments, the OAL uses the "Type 2" OMNI Compressed Header (OCH2) formats shown in Figure 11 and Figure 12:

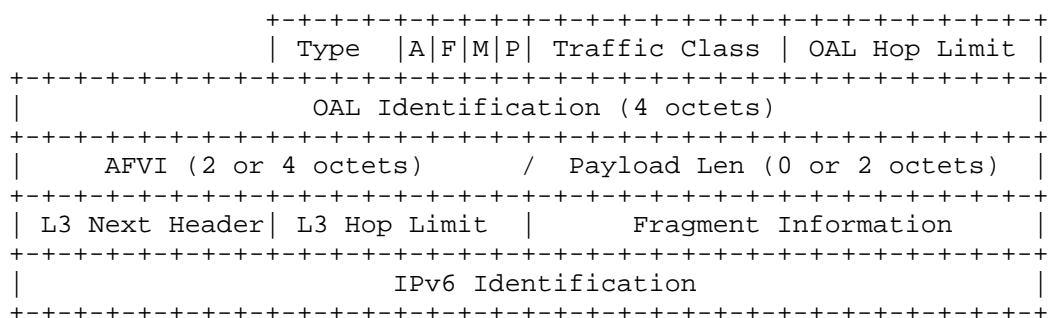


Figure 11: OMNI Compressed Header, Type 2 (OCH2) Format (a)

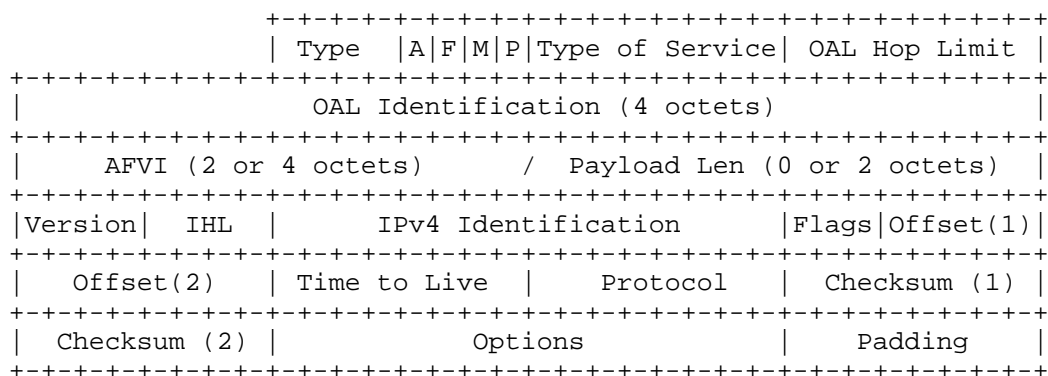


Figure 12: OMNI Compressed Header, Type 2 (OCH2) Format (b)

In both of the above OCH2 formats, the leading octets up to and including the Payload Len (when present) include the same information that would appear in a corresponding OCH1 format (a) header. The (F) flag is set to 0 for OCH2 format (a) or 1 for OCH2 format (b), while all other flags are processed the same as for OCH1 format (a).

The remainder of the OCH2 format (a) includes fields that would appear in an uncompressed IPv6 header per [RFC8200] plus Fragmentation Information. For the standard IPv6 Fragment Header, Fragment Information consists of the 13-bit Fragment Offset followed by the 3 IPv6 Fragment Header flag bits. For the EFH, Fragmentation Information consists of the NH-Cache followed by the 2 EFH flag bits followed by the 6-bit Index.

The remainder of OCH2 format (b) includes fields that would appear in an uncompressed IPv4 header per [RFC0791] with the Options and Padding lengths calculated based on IHL. In both cases, the Source and Destination Addresses are not transmitted.

When an OAL destination or intermediate system receives a carrier packet, it determines the length of the encapsulated OAL information and verifies that the innermost L2 next header field indicates OMNI (see: Section 6.2), then processes any included OMNI L2 extension headers as specified in Section 6.4. The OAL destination then examines the Next Header field of the final L2 extension header. If the Next Header field contains the value TBD1, and the 4-bit Type that follows encodes a value OMNI-IP6, OMNI-OCH1 or OMNI-OCH2 the OAL node processes the remainder of the OAL header as a full or compressed header as specified above.

When an OAL node forwards an OAL packet, it determines the AFVI for the next OAL hop by using the AFVI included in the OCH to search for a matching AFV. The OAL intermediate system then writes the next hop AFVI into the OCH and forwards the OAL packet to the next hop. This same AFVI re-writing progression begins with the OAL source then continues over all OAL intermediate nodes and finally ends at the OAL destination.

If the OAL node is the destination, it instead reconstructs the OAL and original IP headers based on the information cached in the AFV combined with the received information in the OCH1/2. For non-atomic fragments, the OAL node then adds the resulting OAL fragment to the reassembly cache if the Identification is acceptable. Following OAL reassembly if necessary, the OAL node delivers the original IP packet to the network layer.

For all OCH1/2 types, the source node sets all Reserved fields and bits to 0 on transmission and the destination node ignores the values on reception. For both OCH1/2, ECN information is compiled for first fragments, and not for non-first fragments.

Finally, if an IPv6 Hop-by-Hop (HBH) and/or Routing Header extension header is required to appear as per-fragment extensions with each OAL fragment that uses OCH1 format (b) or OCH2 compression the OAL node inserts an OMNI-HBH and/or OMNI-RH header as the first extension(s) following the L2 header and before the OMNI-OCH1/2 as discussed in Section 6.4.

6.6. L2 UDP/IP Encapsulation Avoidance

When the OAL node is unable to determine whether the next OAL hop is connected to the same underlay link, it should perform carrier packet L2 encapsulation for initial packets sent via the next hop over a specific underlay interface by including full UDP/IP headers and with the UDP port numbers set as discussed in Section 6.2. The node can thereafter attempt to send an IPv6 ND solicitation message to the next OAL hop in carrier packet(s) that omit the UDP header and set the IP protocol number to TBD1. If the OAL node receives an IPv6 ND reply, it can omit the UDP header in subsequent packets. The node can further attempt to send an IPv6 ND solicitation in carrier packet(s) that omit both the UDP and IP headers and set EtherType to TBD2. If the source receives an IPv6 ND reply, it can begin omitting both the UDP and IP headers in subsequent packets.

Note: in the above, "next OAL hop" refers to the first OAL node encountered on the optimized path to the destination over a specific underlay interface as determined through route optimization (e.g., see: [I-D.templin-6man-aero3]). The next OAL hop could be a Proxy/Server, Gateway or the OAL destination itself.

6.7. OAL Identification Window Maintenance

The OAL encapsulates each original IP packet as an OAL packet then performs fragmentation to produce one or more carrier packets with the same 8-octet Identification value. In environments where spoofing is not considered a threat, OMNI interfaces send OAL packets with Identifications beginning with an unpredictable Initial Send Sequence (ISS) value [RFC7739] monotonically incremented (modulo $2^{**}64$) for each successive OAL packet sent to either a specific neighbor or to any neighbor. (The OMNI interface may later change to a new unpredictable ISS value as long as the Identifications are assured unique within a timeframe that would prevent the fragments of a first OAL packet from becoming associated with the reassembly of a second OAL packet.) In other environments, OMNI interfaces should

maintain explicit per-flow send and receive windows to detect and exclude spurious carrier packets that might clutter the reassembly cache as discussed below.

OMNI interface neighbors use a window synchronization service similar to TCP [RFC9293] to maintain unpredictable ISS values incremented (modulo 2^{64}) for each successive OAL packet and re-negotiate windows often enough to maintain an unpredictable profile. OMNI interface neighbors exchange IPv6 ND messages that include OMNI Neighbor Synchronization sub-options that include TCP-like information fields and flags to manage streams of OAL packets instead of streams of octets. As a link layer service, the OAL provides low-persistence best-effort retransmission with no mitigations for duplication, reordering or deterministic delivery. Since the service model is best-effort and only control message sequence numbers are acknowledged, OAL nodes can select unpredictable new initial sequence numbers outside of the current window without delaying for the Maximum Segment Lifetime (MSL).

OMNI interface end systems and intermediate systems maintain current and previous per-flow window state in IPv6 ND NCEs and/or AFVs to support dynamic rollover to a new window while still sending OAL packets and accepting carrier packets from the previous windows. OMNI interface neighbors synchronize windows through asymmetric and/or symmetric IPv6 ND message exchanges. When OMNI end and intermediate systems receive an IPv6 ND message with new per-flow window information, it resets the previous window state based on the current window then resets the current window based on new and/or pending information.

The IPv6 ND message OMNI option Neighbor Synchronization sub-option includes TCP-like information fields including Sequence Number, Acknowledgement Number, Window and flags (see: Section 10). OMNI interface neighbors and intermediate systems maintain the following TCP-like state variables on a per-interface-pair basis (i.e., through a combination of NCE and/or AFV state):

Send Sequence Variables (current, previous and pending)

SND.NXT - send next
SND.WND - send window
ISS - initial send sequence number

Receive Sequence Variables (current and previous)

RCV.NXT - receive next
RCV.WND - receive window
IRS - initial receive sequence number

OMNI interface neighbors "OAL A" and "OAL B" exchange IPv6 ND messages per [RFC4861] with OMNI options that include TCP-like information fields in a Neighbor Synchronization. When OAL A synchronizes with OAL B, it maintains both a current and previous SND.WND beginning with a new unpredictable ISS and monotonically increments SND.NXT for each successive OAL packet transmission. OAL A initiates synchronization by including the new ISS in the Sequence Number of an authentic IPv6 ND message with the SYN flag set and with Window set to M (up to 2^{24}) as its advertised send window size while creating a NCE in the INCOMPLETE state if necessary. OAL A caches the new ISS as pending, uses the new ISS as the Identification for OAL encapsulation, then sends the resulting OAL packet to OAL B and waits up to RetransTimer milliseconds to receive an IPv6 ND message response with the ACK flag set (retransmitting up to MAX_UNICAST_SOLICIT times if necessary).

When OAL B receives the SYN, it creates a NCE in the STALE state and also an AFV if necessary, resets its RCV variables and caches the source's send window size M as its receive window size. OAL B then prepares an IPv6 ND message with the ACK flag set, with the Acknowledgement Number set to OAL A's next sequence number, and with Window set to M. Since OAL B does not assert an ISS of its own, it uses the IRS it has cached for OAL A as the Identification for OAL encapsulation then sends the ACK to OAL A.

When OAL A receives the ACK, it notes that the Identification in the OAL header matches its pending ISS. OAL A then sets the NCE state to REACHABLE and resets its SND variables based on the Window size and Acknowledgement Number (which must include the sequence number following the pending ISS). OAL A can then begin sending OAL packets to OAL B with Identification values within the (new) current SND.WND for this interface pair for up to ReachableTime milliseconds or until the NCE is updated by a new IPv6 ND message exchange. This implies that OAL A must send a new SYN before sending more than N OAL packets within the current SND.WND, i.e., even if ReachableTime is not nearing expiration. After OAL B returns the ACK, it accepts carrier packets received from OAL A via this interface pair within either the current or previous RCV.WND as well as any new authentic IPv6 ND messages with the SYN flag set received from OAL A even if outside the windows.

OMNI interface neighbors can employ asymmetric window synchronization as described above using 2 independent (SYN -> ACK) exchanges (i.e., a 4-message exchange), or they can employ symmetric window synchronization using a modified version of the TCP "3-way handshake" as follows:

- * OAL A prepares a SYN with an unpredictable ISS not within the current SND.WND and with Window set to M as its advertised send window size. OAL A caches the new ISS and Window size as pending information, uses the pending ISS as the Identification for OAL encapsulation, then sends the resulting OAL packet to OAL B and waits up to RetransTimer milliseconds to receive an ACK response (retransmitting up to MAX_UNICAST_SOLICIT times if necessary).
- * OAL B receives the SYN, then resets its RCV variables based on the Sequence Number while caching OAL A's send window size M as its receive window size. OAL B then selects a new unpredictable ISS outside of its current window, then prepares a response with Sequence Number set to the pending ISS and Acknowledgement Number set to OAL A's next sequence number. OAL B then sets both the SYN and ACK flags, sets Window to a chosen send window size N and sets the OPT flag according to whether an explicit concluding ACK is optional or mandatory. OAL B then uses the pending ISS as the Identification for OAL encapsulation, sends the resulting OAL packet to OAL A and waits up to RetransTimer milliseconds to receive an acknowledgement (retransmitting up to MAX_UNICAST_SOLICIT times if necessary).
- * OAL A receives the SYN/ACK, then resets its SND variables based on the Acknowledgement Number (which must include the sequence number following the pending ISS). OAL A then resets its RCV variables based on the Sequence Number and OAL B's advertised send Window N and marks the NCE as REACHABLE. If the OPT flag is clear, OAL A next prepares an immediate unsolicited IPv6 ND control message with the ACK flag set, the Acknowledgement Number set to OAL B's next sequence number, with Window set to N, and with the OAL encapsulation Identification to SND.NXT, then sends the resulting OAL packet to OAL B. If the OPT flag is set and OAL A has OAL packets queued to send to OAL B, it can optionally begin sending their carrier packets under the current SND.WND as implicit acknowledgements instead of returning an explicit ACK.
- * OAL B receives the implicit/explicit acknowledgement(s) then resets its SND state based on the pending/advertised values and marks the NCE as REACHABLE. Note that OAL B sets the OPT flag in the SYN/ACK to assert that it will interpret timely receipt of carrier packets within the (new) current window as an implicit acknowledgement. Potential benefits include reduced delays and control message overhead, but use case analysis is outside the scope of this specification.)

Following synchronization, OAL A and OAL B hold updated NCEs and AFVs, and can exchange OAL packets with Identifications set to SND.NXT for each flow while the state remains REACHABLE and there is

available window capacity. (Intermediate systems that establish AFVs for the per-flow window synchronization exchanges can also use the Identification window for source validation.) Either neighbor may at any time send a new SYN to assert a new ISS. For example, if OAL A's current SND.WND for OAL B is nearing exhaustion and/or ReachableTime is nearing expiration, OAL A can continue sending OAL packets under the current SND.WND while also sending a SYN with a new unpredictable ISS. When OAL B receives the SYN, it resets its RCV variables and may optionally return either an asymmetric ACK or a symmetric SYN/ACK to also assert a new ISS. While sending SYNs, both neighbors continue to send OAL packets with Identifications set to the current SND.NXT for each interface pair then reset the SND variables after an acknowledgement is received.

While the optimal symmetric exchange is efficient, anomalous conditions such as receipt of old duplicate SYNs can cause confusion for the algorithm as discussed in Section 3.5 of [RFC9293]. For this reason, the OMNI Neighbor Synchronization sub-option includes an RST flag which OAL nodes set in solicited IPv6 ND message responses to ACKs received with incorrect acknowledgement numbers. The RST procedures (and subsequent synchronization recovery) are conducted exactly as specified in [RFC9293].

OMNI interfaces that employ the window synchronization procedures described above observe the following requirements:

- * OMNI interfaces MUST select new unpredictable ISS values that are at least a full window outside of the current SND.WND.
- * OMNI interfaces MUST set the Window field in SYN messages as a non-negotiable advertised send window size.
- * OMNI interfaces MUST send IPv6 ND messages used for window synchronization securely while using unpredictable initial Identification values until synchronization is complete.

It is essential to understand that the above window synchronization operations between nodes OAL(A) and OAL(B) are conducted in IPv6 ND message exchanges over multihop paths with potentially many OAL(i) intermediate hops in the forward and reverse paths (which may be disjoint). Each such forward path OAL(i) caches the Sequence Number and Window size advertised from OAL(A) to OAL(B) in its AFV entry indexed by the previous hop L2ADDR and AFVI, while each such reverse path OAL(i) caches the Sequence Number, window size and AFVI advertised from OAL(B) to OAL(A). (The forward/reverse path OAL(i) nodes then select new unique next-hop AFVIs before forwarding.)

While multiple independent paths may exist between nodes OAL(A) and OAL(B), the synchronized Sequence Numbers between the two nodes apply collectively to all paths. Nodes OAL(A) and OAL(B) therefore perform initial synchronization through IPv6 ND message exchanges with the SYN flag set over a first path for which intermediate nodes cache the Sequence Number and Window size in their AFVs. However, IPv6 ND message exchanges that establish and maintain alternate paths include the current Sequence Number and residual Window size but with the SYN flag clear.

Each neighbor pair can therefore dynamically coordinate multiple independent paths from a single Sequence Number space in this way. When nodes OAL(A) and OAL(B) need to re-synchronize they again advertise new Sequence Number and Window size values with the SYN flag set. The nodes must then exchange additional IPv6 ND messages using the new values and with the SYN flag clear to establish or maintain alternate paths.

Note: Although OMNI interfaces employ TCP-like window synchronization and support ACK responses to SYNs, all other aspects of the IPv6 ND protocol (e.g., control message exchanges, NCE state management, timers, retransmission limits, etc.) are honored exactly per [RFC4861]. OMNI interfaces further manage per-interface-pair window synchronization parameters in one or more AFVs for each neighbor pair.

Note: Recipients of OAL-encapsulated IPv6 ND messages index the NCE based on the message Source Address, which also determines the carrier packet Identification window. However, IPv6 ND messages may contain a message Source Address that does not match the OMNI encapsulation Source Address when the recipient acts as a proxy.

Note: OMNI interface neighbors apply separate send and receive windows for all of their (multilink) underlay interface pairs that exchange carrier packets. Each interface pair represents a distinct underlay network path, and the set of paths traversed may be highly diverse when multiple interface pairs are used. OMNI intermediate systems therefore become aware of each distinct set of interface pair window synchronization parameters based on periodic IPv6 ND message updates to their respective AFVs.

6.8. OAL Fragmentation Reports and Retransmissions

When the OAL destination experiences reassembly congestion for a specific flow (e.g., when excessive numbers of reassembly failures are occurring), it can send an OAL Fragmentation Report (FRAGREP) message to the OAL source to recommend a reduced Maximum Receive Unit (MRU) for the flow (see: Section 10). When the OAL source received the FRAGREP, it caches the new MRU for the flow and returns "soft errors" to original sources that send larger packets (see: Section 6.9). When the OAL destination experiences reassembly congestion for all flows from the same OAL source, it can return FRAGREP messages with Flow Label set to 0 as indication that all flows are affected.

When the round-trip delay from the original source to the final destination is long while the round-trip time from the OAL source to the OAL destination is significantly shorter, the OAL source can maintain a short-term cache of the OAL fragments it sends to OAL destinations for each flow in case timely best-effort selective retransmission is requested. The OAL destination in turn maintains a checklist for (Source, Destination, Identification)-tuples of recently received OAL fragments and notes the ordinal numbers of OAL fragments already received (i.e., as ordinals #0, #1, #2, #3, etc.). The timeframe for maintaining the OAL source and destination caches determines the link persistence (see: [RFC3366]).

If the OAL destination notices some fragments missing after most other fragments within the same link persistence timeframe have already arrived, it may issue an Automatic Repeat Request (ARQ) with Selective Repeat (SR) by sending an unsolicited IPv6 ND neighbor control message to the OAL source. The OAL destination creates a message with an OMNI option with one or more FRAGREP sub-options that include Bitmaps for fragments received and missing from this OAL source (see: Section 10). The OAL destination includes an authentication signature if necessary, performs OAL encapsulation (with its own address as the OAL Source Address and the Source Address of the message that prompted the unsolicited IPv6 ND as the OAL Destination Address) and sends the message to the OAL source.

If an OAL intermediate system or OAL destination processes an OAL fragment for which corruption is detected, it may similarly issue an immediate ARQ/SR the same as described above. The FRAGREP provides an immediate (rather than time-bounded) indication to the OAL source that a fragment has been lost.

When the OAL source receives the IPv6 ND message, it authenticates the message then examines any enclosed FRAGREPs. For each (Source, Destination, Identification)-tuple, the OAL source determines whether

it still holds the corresponding OAL fragments in its cache and retransmits any for which the Bitmap indicates a loss event. For example, if the Bitmap indicates that ordinal fragments #3, #7, #10 and #13 from the OAL packet with Identification 0x0123456789abcdef are missing the OAL source only retransmits those fragments. When the OAL destination receives the retransmitted OAL fragments, it admits them into the reassembly cache and updates its checklist. If some fragments are still missing, the OAL destination may send a small number of additional IPv6 ND ARQ/SRs within the link persistence timeframe.

The OAL therefore provides a link layer low-to-medium persistence ARQ/SR service consistent with [RFC3366] and Section 8.1 of [RFC3819]. The service provides the benefit of timely best-effort link layer retransmissions which may reduce OAL fragment loss and avoid some unnecessary end-to-end delays. This best-effort network-based service therefore complements transport and higher layer end-to-end protocols responsible for true reliability.

6.9. OMNI Interface MTU Feedback Messaging

When the OMNI interface forwards original IP packets from the network layer, it invokes the OAL and returns internally-generated Path MTU Discovery (PMTUD) ICMPv4 "Fragmentation Needed and Don't Fragment Set" [RFC1191] or ICMPv6 "Packet Too Big (PTB)" [RFC8201] messages as necessary. This document refers to both message types as "PTBs" and introduces a distinction between PTB "hard" and "soft" errors as discussed below.

Ordinary PTB messages are hard errors that always indicate loss due to a real MTU restriction has occurred. However, the OMNI interface can also forward original IP packets via OAL encapsulation and fragmentation while at the same time returning PTB soft error messages (subject to rate limiting) to the original source to suggest smaller sizes due to factors such as link performance characteristics, excessive numbers of fragments needed, reassembly congestion, etc.

This ensures that the path MTU is adaptive and reflects the current path used for a given data flow. The OMNI interface can therefore continuously forward original IP packets without loss while returning PTB soft error messages that recommend smaller sizes. Original sources that receive the soft errors in turn reduce the size of the original IP packets they send the same as for hard errors, but not necessarily due to a loss event. The original source can then resume sending larger packets if the soft errors subside.

OAL intermediate systems that experience fragment loss and OAL end systems that experience reassembly cache congestion can return unsolicited IPv6 ND control messages that include OMNI encapsulated PTB soft error messages to OAL sources that originate fragments (subject to rate limiting). The OAL node creates a secured control message with an OMNI option containing an ICMPv6 Error sub-option. The OAL node encodes a PTB message in the sub-option with MTU set to a reduced value and with the leading portion an OAL first fragment containing the header of an original IP packet for which the source must be notified (see: Section 10).

The OAL node that sends the IPv6 ND message encapsulates the leading portion of the OAL first fragment (beginning with the OAL header) in the PTB "packet in error" field and signs the message if an authentication signature is included. The OAL node then performs OAL encapsulation (with its own address as the Source Address and the Source Address of the message that prompted the IPv6 ND response as the Destination Address) and sends the message to the OAL source. (Note that OAL intermediate systems forward IPv6 ND control messages via the secured spanning tree while OAL source and destination end systems include an authentication signature when necessary.)

The OAL source prepares the PTB soft error by first setting the Type field to 2 for IPv6 [RFC4443] or "Packet Too Big" for IPv4 (see: [I-D.templin-6man-ipid-ext2]). The OAL source then sets the Code field to "PTB Soft Error (no loss)" if the OAL destination forwarded the original IP packet successfully or "PTB Soft Error (loss)" if it was dropped (see: [I-D.templin-6man-ipid-ext2]). The OAL source next sets the PTB Destination Address to the original IP packet Source Address, and sets the PTB Source Address to one of its OMNI interface addresses that is reachable from the perspective of the original source.

The OAL source then sets the MTU field to a value smaller than the original IP packet size but no smaller than 1280, writes as much of the original IP packet first fragment as possible into the "packet in error" field such that the entire PTB including the IP header is no larger than 1280 octets for IPv6 or 576 octets for IPv4. The OAL source then calculates and sets the ICMP Checksum and returns the PTB to the original source.

An original sources that receives these PTB soft errors first verifies that the ICMP Checksum is correct and the packet-in-error contains the leading portion of one of its recent packet transmissions. The original source can then adaptively tune the size of the original IP packets it sends to produce the best possible throughput and latency, with the understanding that these parameters may fluctuate over time due to factors such as congestion, mobility,

network path changes, etc. Original sources should therefore consider receipt or absence of soft errors as hints of when decreasing or increasing packet sizes may provide better performance.

The OMNI interface supports continuous transmission and reception of packets of various sizes in the face of dynamically changing network conditions. Moreover, since PTB soft errors do not indicate a hard limit, original sources that receive soft errors can resume sending larger packets without waiting for the recommended 10 minutes specified for PTB hard errors [RFC1191][RFC8201]. The OMNI interface therefore provides an adaptive service that accommodates MTU diversity especially well-suited for air/land/sea/space mobile Internetworking.

Note: when the OAL source receives persistent Fragmentation Reports for a given flow (see: Section 6.8), it should return PTB soft errors to the original source (subject to rate limiting) the same as if it had received PTB soft errors from the OAL destination. When the original source is likely to retransmit an entire original IP packet on its own behalf in case of loss, the OAL destination can elect to return only PTB soft errors and refrain from returning Fragmentation Reports.

Note: the OAL source may receive control messages that include both a PTB soft error and Fragmentation Report(s). If so, the OAL source both returns PTB soft errors to the original source (subject to rate limiting) and retransmits any missing fragments if it is configured to do so.

6.10. OAL Composite Packets

The OAL source ordinarily includes a 40-octet IPv6 encapsulation header for each original IP packet during OAL encapsulation. The OAL source then performs fragmentation such that a copy of the 40-octet IPv6 header plus a 16-octet IPv6 Extended Fragment Header is included in each OAL fragment (when a Routing Header is added, the OAL encapsulation headers become larger still). However, these encapsulations may represent excessive overhead in some environments.

OAL header compression as discussed in Section 6.5 can dramatically reduce encapsulation overhead, however a complementary technique known as "packing" (see: [I-D.ietf-intarea-tunnels]) supports encapsulation of multiple original IP packets and/or control messages within a single OAL "composite packet".

When the OAL source has multiple original IP packets to send to the same OAL destination with total length no larger than the OAL destination EMTU_R, it can concatenate them into a composite packet

encapsulated in a single OAL header. Within the OAL composite packet, the IP header of the first original IP packet (iHa) followed by its data (iDa) is concatenated immediately following the OAL header. The IP header of the next original packet (iHb) followed by its data (iDb) is then concatenated immediately following the first, with each remaining original IP packet concatenated in succession. The OAL composite packet format is transposed from [I-D.ietf-intarea-tunnels] and shown in Figure 13:

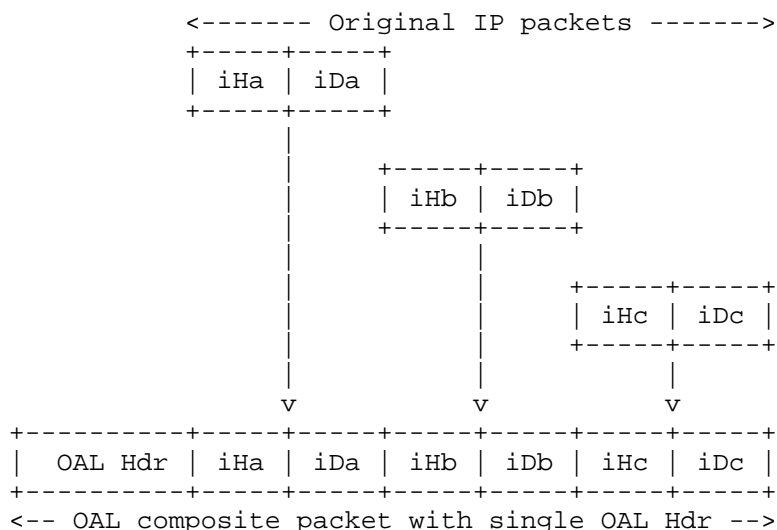


Figure 13: OAL Composite Packet Format

When the OAL source prepares a composite packet, it applies OAL fragmentation then applies L2 encapsulation and sends the resulting carrier packets to the OAL destination. When the OAL destination receives the composite packet it first reassembles if necessary. The OAL destination then selectively extracts each original IP packet (e.g., by setting pointers into the composite packet buffer and maintaining a reference count, by copying each packet into a separate buffer, etc.) and forwards each one to the network layer. During extraction, the OAL determines the IP protocol version of each successive original IP packet 'j' by examining the 4 most-significant bits of $iH(j)$, and determines the length of each one by examining the rest of $iH(j)$ according to the IP protocol version.

When an OAL source prepares a composite packet that includes an IPv6 ND message as the first original IP packet (i.e., iHa/iDa) it includes any additional original IP packets in concatenated succession then includes a trailing OMNI option. If the OMNI option contains an authentication sub-option, the OAL source calculates the

authentication signature over the entire length of the composite packet. (A second common use case entails a path MTU probe beginning with an unsigned IPv6 ND message followed by a suitably large NULL packet (e.g., an IP packet with padding octets added beyond the IP header and with {Protocol, Next Header} set to 59 ("No Next Header"), a UDP/IP packet with port number set to 9 ("discard") [RFC0863], etc.)

The OAL source can also apply this composite packet packing technique at the same time it performs OCH1 header compression as discussed in Section 6.5. Note that this technique can only be applied for original IP packets of a single flow, such as for a stream of packets for the flow that are queued for transmission service at roughly the same time.

6.11. OAL Bubbles

OAL sources may send NULL OAL packets known as "bubbles" for the purpose of establishing Network Address Translator (NAT) state on the path to the OAL destination. The OAL source prepares a bubble by crafting an OAL header with appropriate IPv6 Source and Destination ULAs, with the IPv6 Next Header field set to the value 59 ("No Next Header" - see [RFC8200]) and with 0 or more octets of NULL protocol data immediately following the IPv6 header.

The OAL source includes a random Identification value then encapsulates the OAL packet in L2 headers destined to either the mapped address of the OAL destination's first-hop ingress NAT or the L2 address of the OAL destination itself. When the OAL source sends the resulting carrier packet, any egress NATs in the path toward the L2 destination will establish state based on the activity. At the same time, the bubble themselves will be harmlessly discarded by either an ingress NAT on the path to the OAL destination or by the OAL destination itself.

The bubble concept for establishing NAT state originated in [RFC4380] and was later updated by [RFC6081]. OAL bubbles may be employed by mobility services such as AERO.

6.12. OAL Requirements

In light of the above, OAL sources, destinations and intermediate systems observe the following normative requirements:

- * OAL sources MUST forward original IP packets either larger than the OMNI interface minimum EMTU_R or smaller than the minimum OFS as atomic fragments (i.e., and not as multiple fragments).

- * OAL sources MUST perform OAL fragmentation such that all non-final fragments are equal in length while the final fragment may be a different length.
- * OAL sources MUST produce non-final fragments with payloads no smaller than the minimum OFS during fragmentation.
- * OAL intermediate systems SHOULD and OAL destinations MUST unconditionally drop any non-final OAL fragments with payloads smaller than the minimum OFS.
- * OAL destinations MUST drop any new OAL fragments that would overlap with other fragments and/or leave holes smaller than the minimum OFS between fragments that have already been received.

Note: Certain legacy network hardware of the past millennium was unable to accept IP fragment "bursts" resulting from a fragmentation event - even to the point that the hardware would reset itself when presented with a burst. This does not seem to be a common problem in the modern era, where fragmentation and reassembly can be readily demonstrated at line rate (e.g., using tools such as 'iperf3') even over fast links on ordinary hardware platforms. Even so, while the OAL destination is reporting reassembly congestion (see: Section 6.9) the OAL source could impose "pacing" by inserting an inter-fragment delay and increasing or decreasing the delay according to congestion indications.

6.13. OAL Fragmentation Security Implications

As discussed in Section 3.7 of [RFC8900], there are 4 basic threats concerning IPv6 fragmentation; each of which is addressed by effective mitigations as follows:

1. Overlapping fragment attacks - reassembly of overlapping fragments is forbidden by [RFC8200]; therefore, this threat does not apply to the OAL.
2. Resource exhaustion attacks - this threat is mitigated by providing a sufficiently large OAL reassembly cache and instituting "fast discard" of incomplete reassemblies that may be part of a buffer exhaustion attack. The reassembly cache should be sufficiently large so that a sustained attack does not cause excessive loss of good reassemblies but not so large that (timer-based) data structure management becomes computationally expensive. The cache should also be indexed based on the arrival underlay interface such that congestion experienced over a first underlay interface does not cause discard of incomplete reassemblies for uncongested underlay interfaces.

3. Attacks based on predictable fragment Identification values - in environments where spoofing is possible, this threat is mitigated through the use of Identification windows beginning with unpredictable values per Section 6.7. By maintaining windows of acceptable Identifications, OAL neighbors can quickly discard spurious carrier packets that might otherwise clutter the reassembly cache.
4. Evasion of Network Intrusion Detection Systems (NIDS) - since the OAL source employs a robust OFS, network-based firewalls can inspect and drop OAL fragments containing malicious data thereby disabling reassembly by the OAL destination. However, each OAL destination should also employ a (host-based) firewall.

IPv4 includes a 2-octet (16-bit) Identification (IP ID) field with only 65535 unique values such that even at moderate data rates the field could wrap and apply to new carrier packets while the fragments of old carrier packets using the same IP ID are still alive in the network [RFC4963]. However, IPv4 links that configure a small MTU are likely to occur only at extreme network edges where low data rate links occur [RFC3819]. Since IPv6 provides a 4-octet (32-bit) Identification value, IP ID wraparound for IPv6 fragmentation may only be a concern at extreme data rates (e.g., 1Tbps or more). These limitations are fully addressed through the 8-octet (64-bit) Extended Identification format supported by [I-D.templin-6man-ipid-ext2].

Unless the path is secured at the network layer or below (i.e., in environments where spoofing is possible), OMNI interfaces MUST NOT send OAL packets/fragments with Identification values outside the current window and MUST secure IPv6 ND messages used for address resolution or window state synchronization. OAL destinations SHOULD therefore discard without reassembling any out-of-window OAL fragments received over an unsecured path.

6.14. Control/Data Plane Considerations

The above sections primarily concern data plane aspects of the OMNI interface service and describe the data plane service model offered to the network layer. OMNI interfaces also internally employ a control plane service based on IPv6 ND messaging. These control plane messages are first subject to OAL encapsulation then forwarded over secured underlay interfaces (e.g., IPsec tunnels, secured direct point-to-point links, etc.) or over unsecured underlay interfaces and with an authentication signature included.

OMNI interfaces must send all control plane messages as "atomic OAL packets". This means that these messages must not be subjected to OAL fragmentation and reassembly, although they may be subjected to

L2 fragmentation and reassembly along some paths. Fragmentation security concerns for large IPv6 ND messages are documented in [RFC6980].

7. Ethernet-Compatible Link Layer Frame Format

When the OMNI interface forwards original IP packets from the network layer it first invokes OAL encapsulation and fragmentation, then wraps each resulting OAL packet/fragment in any necessary L2 headers to produce carrier packets according to the native frame format of the underlay interface. For example, for Ethernet-compatible interfaces the frame format is specified in [RFC2464], for aeronautical radio interfaces the frame format is specified in standards such as ICAO Doc 9776 (VDL Mode 2 Technical Manual), for various forms of tunnels the frame format is found in the appropriate tunneling specification, etc.

When the OMNI interface encapsulates an OAL packet/fragment directly over an Ethernet-compatible link layer, the over-the-wire transmission format is shown in Figure 14:




Figure 14: OMNI Ethernet Frame Format

The format includes a standard Ethernet Header ("eth-hdr") with EtherType TBD2 (see: Section 21.2) followed by an Ethernet Payload that includes zero or more OMNI Extension Headers followed by an OAL (or native IPv6/IPv4) Packet/Fragment. The Ethernet Payload is then followed by a standard Ethernet Trailer ("eth-trail").

The first OMNI extension header and the OAL Packet/Fragment both begin with a 4-bit "Type/Version" as discussed in Section 6.2. When "Type/Version" encodes an OMNI extension header type, the length of the extension headers is limited by [I-D.ietf-6man-eh-limits] and the length of the OAL Packet/Fragment is determined by the IP header fields that follow the extension headers.

When "Type/Version" encodes OMNI-OCH1/2, OMNI-IP4 or OMNI-IP6 the length of the OAL Packet/Fragment is determined by the {Total, Payload} Length field found in the full/compressed header according to the specific protocol rules.

See Figure 2 for a map of the various L2 layering combinations possible. For any layering combination, the final layer (e.g., UDP, IP, Ethernet, etc.) must have an assigned number and frame format representation that is compatible with the selected underlay interface.

8. OMNI Addressing

OMNI addressing observes the IPv6 addressing architecture [RFC4291] requirement that: "IPv6 addresses of all types are assigned to interfaces, not nodes. An IPv6 unicast address refers to a single interface. Since each interface belongs to a single node, any of that node's interfaces' unicast addresses may be used as an identifier for the node." OMNI addressing further follows the IPv6 address selection policies specified in [RFC6724] as updated by [I-D.ietf-6man-rfc6724-update].

Each OMNI interface is configured over a set of underlay interfaces as a virtual data link layer for the OAL. OMNI nodes assign IP addresses to their underlay interfaces according to the native *NET autoconfiguration service(s) or through manual configuration. OMNI nodes assign IPv6 addresses to their OMNI interfaces as specified in this section.

[RFC4861] requires that hosts and routers assign Link-Local Addresses (LLAs) to all interfaces including the OMNI interface, and that routers use their LLAs as the Source Address for RA and Redirect messages. The OMNI interface satisfies this property by maintaining an internal mapping cache to present the network layer with an LLA-based view of all neighbors while the adaptation layer within the OMNI interface maps IPv6 message Source and Destination LLAs to Multilink Local Addresses (MLAs). (If the node assigns multiple LLAs to the OMNI interface, e.g., as suggested by [I-D.link-6man-gulla] it must also assign multiple MLAs in 1x1 correspondence. In that case, the node would appear as multiple separate nodes on the OMNI link.)

[I-D.templin-6man-mla] specifies MLA types that OMNI nodes can assign to the OMNI interface given sufficient uniqueness and authentication assurances. Candidate MLA types include the Host Identity Tag (HIT) [RFC7343], Hierarchical HIT (HHIT) [RFC9374], and Segment Routing over IPv6 (SRv6) Segment Identifiers [RFC9602] but could also include future special-purpose IPv6 address types identified by the IPv6 prefix. The node assigns an MLA to an OMNI interface configured over its set of underlay interfaces per the IPv6 scoped addressing architecture "site" abstraction [RFC4007]. MLAs are considered as adaptation layer addresses in the architecture.

When the data link layer presents an OAL-encapsulated IPv6 packet with MLA Source/Destination Addresses to the OMNI interface, the adaptation layer decapsulates the IPv6 packet if the MLA Destination Address matches its own address then rewrites the Destination Address with its own LLA while forwarding the packet to the network layer. For IPv6 ND messages that install the neighbor's LLA in the neighbor cache, the adaptation layer first generates a new local LLA for the neighbor with a randomized Modified EUI-64 format interface identifier per [RFC4291] that is unique among all current neighbor LLAs. For all IPv6 ND messages that include a Source/Target Link Layer Address Option (S/TLLAO) the adaptation layer then rewrites the S/TLLAO based on the EUI-64 format interface identifier from the local LLA. For all IPv6 packets with MLA Source Addresses, the adaptation layer then rewrites the Source Address with the local LLA before presenting the IPv6 packet to the network layer.

When the network layer presents an IP packet with LLA Source/Destination Addresses to an OMNI interface, the adaptation layer rewrites the LLA Source Address with its own MLA and rewrites the LLA Destination Address to the MLA of the peer or to a site-scoped multicast or anycast address. The OMNI interface then encapsulates the packet in an OAL header with MLA Source/Destination Addresses and presents it to the data link layer. The adaptation layer mapping table therefore must ensure that the network layer sees a unique local representation of the LLA for each active neighbor while mapping its local S/TLLAO view to the neighbor's view and while including only MLAs (and not LLAs) in actual message exchanges with neighbors.

OMNI interfaces assign IPv6 Unique Local Addresses (ULAs) and use them as the Source and Destination Addresses in IPv6 packets forwarded over the OMNI interface within the local OMNI link segment. OMNI interfaces also assign corresponding IPv6 Globally Unique Addresses (GUAs) and use them as Source and Destination Addresses for IPv6 packets exchanged with peers in external networks. ULAs are routable only within the scope of an OMNI link segment, and are derived from the IPv6 prefix fd00::/8 (i.e., the ULA prefix fc00::/7 followed by the L bit set to 1). The 56 bits following fd00::/8 encode a 40-bit Global ID followed by a 16-bit Subnet ID followed by a 64-bit Interface Identifier as specified in Section 3 of [RFC4193].

When a Proxy/Server configures a ULA prefix for OMNI, it selects a 40-bit Global ID for the OMNI link segment initialized to a candidate pseudo-random value as specified in Section 3 of [RFC4193]. All nodes on the same OMNI link segment use the same Global ID, and statistical uniqueness of the pseudo-random Global ID provides a unique OMNI link segment identifier. This property allows different link segments to join together in the future without requiring renumbering even if the segments come in contact with one another and overlap, e.g., as a result of a mobility event.

Proxy/Servers for each OMNI link segment use the DHCPv6 service to delegate 1x1 mapped ULA/GUA SNP addresses for each Client that requests an address delegation. (A suitable method for SNP GUA address delegation appears in [I-D.gont-dhcgw-dhcpv6-iids], while the corresponding ULA is formed by appending the 64-bit GUA interface identifier to the 64-bit ULA prefix.) Clients in turn assign the ULA/GUA delegations to their OMNI interfaces which ensures that the addresses are available for use and that no duplicates will be assigned within each OMNI link segment. IPv6 address selection at the network layer above the OMNI interface then determines the underlay interface used for service at the data link layer below the OMNI interface. Further considerations for 1x1 ULA/GUA address mapping are discussed in [I-D.ietf-v6ops-ula-usage-considerations] and [RFC6296].

The OMNI link extends across one or more underlying Internetworks to include all Proxy/Servers and other service nodes. All Clients are also considered to be connected to the OMNI link, however unnecessary encapsulations are omitted whenever possible to conserve bandwidth (see: Section 12). An OMNI domain consists of one or more OMNI links joined together to provide service for a common set of MSPs.

OMNI domains include one or more OMNI links that together coordinate a common set of MSPs delegated from the IP GUA prefix space [RFC4291] from which the MS delegates MNPs to support Client PI addressing. OMNI Proxy Servers also configure an SNP paired with a ULA prefix configured as above to delegate PA internal (ULA) and external (GUA) addresses to Clients within their local *NETs.

For IPv6, MSPs are assigned to an OMNI domain by IANA and/or an associated Regional Internet Registry [IPV6] such that the link(s) can be connected to the global IPv6 Internet without causing routing inconsistencies. Instead of GUAs, an OMNI link could use ULAs with the 'L' bit set to 0 (i.e., from the "ULA-C" prefix fc00::/8) [RFC4193], however this would require IPv6 NAT if the domain were ever connected to the global IPv6 Internet.

For IPv4, MSPs are assigned to an OMNI domain by IANA and/or an associated RIR [IPV4] such that the link(s) can be connected to the global IPv4 Internet without causing routing inconsistencies. An OMNI *NET could instead use private IPv4 prefixes (e.g., 10.0.0.0/8, etc.) [RFC6890], however this would require IPv4 NAT at the *NET boundary. OMNI interfaces advertise IPv4 MSPs into IPv6 routing systems as "6to4 prefixes" [RFC3056] (e.g., the IPv6 prefix for the IPv4 MSP "V4ADDR/24" is 2002:V4ADDR::/40).

IPv4 routers that configure OMNI interfaces advertise the prefix TBD3/N (see: IANA Considerations) into the routing systems of their connected *NETs and assign the IPv4 OMNI anycast address TBD3.1 to their *NET interfaces. IPv6 routers that configure OMNI interfaces advertise the prefix 2002:TBD3::/(N+16) into the routing systems of their connected *NETs and assign the IPv6 OMNI anycast address 2002:TBD3:: to their *NET interfaces.

Proxy/Server OMNI interfaces configure ULA/GUA IPv6 SNP SRA addresses per [RFC4291] and accept packets addressed to the SRA the same as for any IPv6 router. Proxy/Servers also configure the global IPv6 SRA address for each MSP managed by this OMNI link and accept packets addressed to the SRA address on their internal interfaces to support Client OMNI link discovery. Client OMNI interfaces configure the IPv6 SRA address corresponding to their MNP delegations.

OMNI interfaces use their OMNI IPv6 and IPv4 anycast addresses to support control plane Service Discovery in the spirit of [RFC7094], i.e., the addresses are not intended for use in supporting longer term data plane flows. Specific applications for OMNI IPv6 and IPv4 anycast addresses are discussed throughout the document as well as in [I-D.templin-6man-aero3].

OMNI Clients and Proxy/Servers use their MLAs as OAL Source and Destination Addresses within the FHS *NET. FHS Proxy/Servers rewrite OAL Source and Destination MLAs as SNP GUAs before forwarding packets over intervening Internetworks on the paths to LHS Proxy/Servers. LHS Proxy servers in turn rewrite OAL Source and Destination SNP GUAs as MLAs for forwarding within the LHS *NET. (Proxy/Servers rewrite the OAL Source Address in the same manner as for NATs and rewrite the OAL Destination Address based on information found in an included SRH.)

9. Node Identification

OMNI Clients and Proxy/Servers that connect over open Internetworks include a unique node identification value for themselves in the IPv6 Source Address and/or in an OMNI option of their IPv6 ND messages (see: Section 10.2.3). Each node configures and includes an MLA as a node identification as discussed in Section 8. (The Universally Unique Identifier (UUID) [RFC9562] is another example of a node identifier which can be self-generated by a node without supporting infrastructure with very low probability of collision.)

When a Client is truly outside the context of any infrastructure, it may have no addressing information at all. In that case, the Client can use an MLA as an IPv6 Source/Destination Address for sustained communications in Vehicle-to-Vehicle (V2V) and (multihop) Vehicle-to-Infrastructure (V2I) scenarios. The Client can also propagate the MLA into the multihop routing tables of (collective) Mobile/Vehicular Ad-hoc Networks (MANETs/VANETs) using only the vehicles themselves as communications relays. MLAs provide an especially useful node identification construct since they appear as properly-formed IPv6 addresses.

When a Client already includes its MLA in the IPv6 Source Address of an original IP packet or IPv6 ND message, it need not also include the MLA in an OMNI Node Identification sub-option.

10. Address Mapping - Unicast

OMNI interfaces maintain network layer conceptual Neighbor and Destination Caches per [RFC1256][RFC4861] the same as for any IP interface. The network layer maintains state through static and/or dynamic Neighbor/Destination Cache Entry (NCE/DCE) configurations.

Each OMNI interface also maintains an internal adaptation layer view of the neighbor cache that supplements the network layer NCEs for each of its active neighbors. For each peer NCE, neighbors also maintain AERO Flow Vectors (AFVs) in the OAL which map per-interface-pair parameters. Throughout this document, the terms "neighbor cache", "NCE" and "AFV" refer to this OAL neighbor cache view unless otherwise specified.

When a Client's network layer sends or receives IPv6 Neighbor Discovery (ND) messages over an OMNI interface, it follows the procedures in [RFC4861] using the Source/Target Link-Layer Address Option (S/TLLAO) format defined for Ethernet [RFC2464]. On transmission, the OMNI interface OAL leaves the S/TLLAO unchanged. On reception, the OAL uses the IPv6 Source Address to translate the S/TLLAO Ethernet address into a unique locally-generated value for this neighbor.

When a Client's network layer sends or receives an ordinary IP packet over an OMNI interface, the OAL consults the Ethernet to OAL IPv6 address mappings established by earlier IPv6 ND message exchanges. On transmission, the OAL uses the Ethernet destination address to determine the Destination Address for an OAL encapsulation header while including an SRH extension if necessary. On reception, the OAL uses the IPv6 encapsulation header Source Address to determine the source address for the virtual Ethernet header.

The OMNI interface must therefore maintain internal per neighbor NCEs that map local Ethernet addresses to remote Ethernet addresses and GUAs/MLAs while exposing only the local representation of the addresses to the IP layer. When the OMNI interface discovers a new neighbor (e.g., when it creates a new NCE based on receipt of an IPv6 ND message), it maps the remote Ethernet address and GUA/MLA to a randomly-chosen 6 octet local Ethernet address that must be unique for this interface then installs the mapping in the cache. When the OMNI interface discards an existing neighbor (e.g., when it deletes an expired NCE), it removes the internal address mappings from the cache.

When the OAL forwards IPv6 ND messages from the network layer to the link layer, it performs encapsulation by adding an adaptation layer IPv6 header (plus any necessary routing headers) and a new pseudo IPv6 ND option trailer that encodes OMNI link-specific information. When the OAL forwards IPv6 ND messages from the link layer to the network layer, it performs decapsulation by removing the adaptation layer IPv6 header while also parsing and removing the trailer. Hence, this document defines a new pseudo IPv6 ND option type termed the "OMNI option" designed for these purposes. Since the pseudo-option is inserted and removed by the adaptation layer and never examined by the network layer, it does not require a formal IPv6 ND option number assignment.

OMNI interface Clients such as aircraft typically have multiple wireless data link types (e.g. satellite-based, cellular, terrestrial, air-to-air directional, etc.) with diverse performance, cost and availability properties. The OMNI interface would therefore appear to have multiple L2 connections, and may include information

for multiple underlay interfaces in a single IPv6 ND message exchange. OMNI interfaces manage their dynamically-changing multilink profiles by including OMNI options in IPv6 ND messages as discussed in the following subsections.

10.1. The OMNI Option

During OAL IPv6 encapsulation of each IPv6 ND message, the OAL source appends a single OMNI (pseudo-)option as a contiguous block of data immediately following the end of the (composite) packet and includes the length of the option in the OAL IPv6 header Payload Length.

During decapsulation of each IPv6 ND message, the OAL destination processes the OMNI option contents then removes the option before delivering the original IPv6 ND message (plus any additional original IP packets from the composite packet) to the network layer.

The OMNI option therefore appears if and only if the (composite) packet begins with an IPv6 ND message. The OMNI option format is shown in Figure 15.

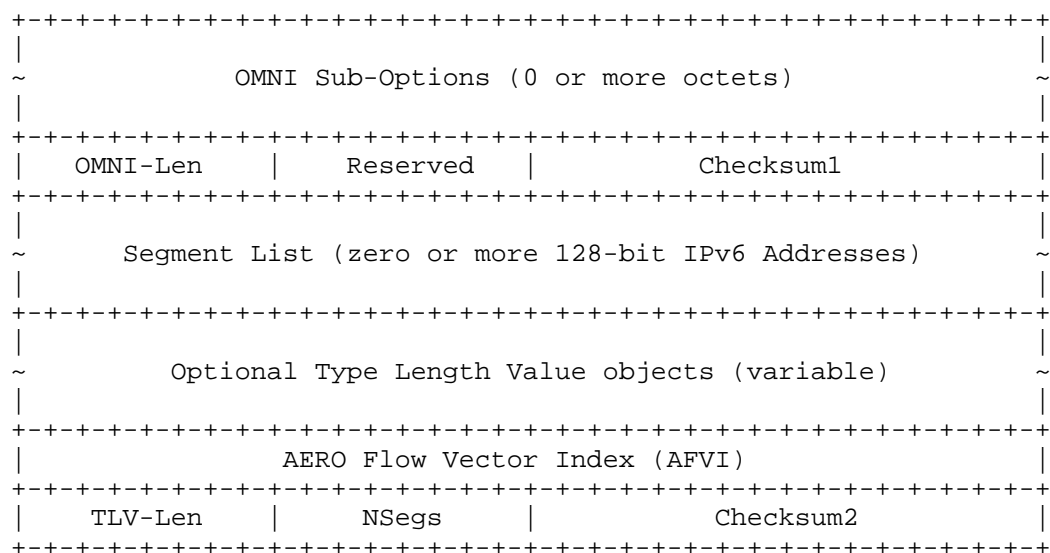


Figure 15: OMNI Option Format

In this format:

- * OMNI Sub-Options is a variable-length concatenation of 0 or more sub-option entries formatted as specified in Section 10.2 such that the total length of all Sub-Options is an integer multiple of

8 octets long, i.e., even if padding octets are necessary. The final sub-option is followed by a variable-length trailer containing the fields described below.

- * OMNI-Len is an 8-bit unsigned integer that encodes the combined integer length of all Sub-Options in 8-octet units. If there are no OMNI sub-options, OMNI-Len encodes the value 0.
- * Reserved is a 1-octet reserved field, set to 0 on transmission and ignored on reception.
- * Checksum1 is a 2-octet Internet checksum calculated over the length of the OMNI option beginning with the first Sub-Options octet and extending to include the OMNI-Len and Reserved fields. The OAL source calculates the Internet checksum and writes the resulting value into the Checksum1 field. The OAL destination verifies the checksum upon receipt and processes the OMNI option further only if the checksum is correct. OAL intermediate systems do not verify the OMNI option checksum and simply pass the option contents up to and including the Checksum1 field unchanged.
- * Segment List records the MLAs of endpoint OAL intermediate systems on the path from the original source Client to its FHS Proxy/Server. For RS messages only, each successive endpoint intermediate system "Proxy/Client" appends its MLA as the final IPv6 address in the Segment List then increments the OAL IPv6 Payload Length by 16 and increments NSegs by 1 (see below).
- * Optional Type Length Value (TLV) objects are included the same as specified in [RFC8754] and may include a Hashed Message Authentication Code (HMAC) [RFC2104] which covers the AFVI and Segment List only. When included, the HMAC is checked then reset by each successive OAL intermediate system in a hop-by-hop fashion. If resetting the HMAC causes its length to change, the OAL intermediate system must also reset both TLV-Len and the OAL IPv6 Payload Length accordingly.
- * AERO Flow Vector Index (AFVI) is a 4-octet field initialized by the IPv6 ND message source for each independent flow and rewritten by each transit and endpoint OAL intermediate system on the path ending at the IPv6 ND message destination (the special value 0 denotes "AFVI unspecified"). The OAL source can then begin sending OAL packets with OCH headers that include the AFVI which each forwarding OAL intermediate system in the path can use to determine the next OAL hop.

- * TLV-Len is the length in octets of the TLV objects field, and provides an offset to the beginning of the TLV objects (or the value 0 if the TLV objects field is null).
- * NSegs includes the number of IPv6 addresses in the Segment List, initialized to 0 by the OAL source for all IPv6 ND message types. For RS messages only, each successive endpoint intermediate system updates the Segment List and OAL IPv6 Payload Length (see above) then increments NSegs by 1.
- * Checksum2 is the Internet checksum calculated beginning with the OAL IPv6 pseudo-header per Section 8.1 of [RFC8200] then continuing over the AFVI, Segment List, TLV objects then finally the TLV-Len and NSegs fields. The OAL source calculates the Internet checksum and writes the resulting value into the Checksum2 field. Each OAL intermediate system up to and including the OAL destination verifies the checksum upon receipt and processes the OMNI option further only if the checksum is correct. The intermediate system then adjusts the OAL header and trailer fields as necessary and resets Checksum2 before forwarding to the next hop.

OMNI encapsulated IPv6 ND messages exchanged over unsecured *NETs between peer Clients or Clients and their Proxy/Servers use either SEND per [RFC3971] or HMAC per [RFC8754][RFC2104] as an adaptation layer authentication service. Since the adaptation layer already applies authentication from within the OMNI interface, the network layer need not also apply IPv6 ND message authentication over the OMNI interface unless there is some reason to propagate a digital signature to the final destination. The OMNI option therefore provides sub-options to support either SEND or HMAC as adaptation layer authentication services.

Although originally specified to operate with Cryptographically Generated Addresses (CGAs) per [RFC3972], SEND notes that: "This specification also allows a node to use non-CGAs with certificates that authorize their use. However, the details of such use are beyond the scope of this specification and are left for future work." Since CGAs do not support verification through an address registration and certification service, OMNI specifically requires alternative MLA types that can satisfy these properties.

The OMNI Sub-Options may include full or partial information for the neighbor. The OMNI interface therefore retains the union of the most recently received information in the corresponding NCE.

10.2. OMNI Sub-Options

The OMNI option includes a Sub-Options block containing zero or more individual sub-options in standard Type-Length-Value (TLV) format. Each successive sub-option is concatenated immediately following its predecessor. All sub-options are encoded as follows:

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| Sub-Type | Sub-Length | Sub-Option Data ...
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 16: Sub-Option Format

- * Sub-Type is a 8-bit field that encodes the sub-option type. Sub-option types defined in this document include:

Sub-Option Name	Sub-Type
Pad1	0
PadN	1
Node Identification	2
CGA	3
RSA Signature	4
Timestamp	5
Nonce	6
Trust Anchor	7
Certificate	8
HMAC	9
Neighbor Synchronization	10
Interface Attributes	11
Traffic Selector	12
Geo Coordinates	13
DHCPv6 Message	14
PIM-SM Message	15
Fragmentation Report	16
ICMPv6 Error	17
Proxy/Server Departure	18

Figure 17

Unassigned Sub-Types are available for future assignment, except that Sub-Types 253 and 254 are reserved for experimentation while Sub-Type 255 is reserved by IANA.

- * Sub-Length is an 8-bit field that encodes the length of the Sub-Option Data in octets (i.e., not including the Sub-Type and Sub-Length fields themselves).

- * Sub-Option Data is a block of data with format determined by Sub-Type and length determined by Sub-Length. Note that each sub-option is concatenated immediately following the previous and may therefore begin and/or end on an arbitrary octet boundary.

The OMNI interface codes all sub-options in a single OMNI option in the same IPv6 ND message in the intended order of processing. If the size of the sub-options would cause the IPv6 ND message to exceed the path MTU, the OMNI interface includes as many sub-options as possible and codes any remaining sub-options in additional IPv6 ND messages.

Various option alignment guidelines appear in [RFC4861], [RFC8200] and [I-D.ietf-dhc-rfc8415bis]. The former two standards ensure that IPv6 options are aligned on natural octet boundaries for 1, 2, 4 and 8-octet fields while the latter permits DHCPv6 options to appear on arbitrary octet boundaries. When the source appends an OMNI option, it must observe the [RFC8200], Appendix A alignment guidelines to ensure that field lengths up to 8 octets are aligned on natural boundaries for the following sub-options: Neighbor Synchronization, Interface Attributes, Traffic Selector, Fragmentation Report and Proxy/Server Departure. Unless otherwise specified, the source may align all other sub-options on arbitrary octet boundaries.

The OMNI interface processes the OMNI option received in an IPv6 ND message while skipping over and ignoring any unrecognized sub-options. If an individual sub-option length would cause processing to exceed the OMNI option instance and/or IPv6 ND message lengths, the OMNI interface accepts any sub-options already processed and ignores the remainder of that instance.

IPv6 ND messages that require OMNI authentication services include an RSA Signature or HMAC sub-option as the first sub-option. A single IPv6 ND message includes a single effective OMNI authentication service sub-option; if multiple are included, the first sub-option is processed and all others are ignored.

Note: large objects that exceed the maximum Sub-Option Data length are not supported under the current specification; if this proves to be limiting in practice, future specifications may define support for fragmenting large sub-options across multiple IPv6 ND messages, if necessary.

The following sub-option types and formats are defined in this document:

10.2.1. Pad1

```

+-----+
| Sub-Type=0 |
+-----+

```

Figure 18: Pad1

- * Sub-Type is set to 0. If multiple consecutive instances of Pad1 appear in the OMNI option of the same message the message should be dropped. (If more than a single consecutive octet of padding is necessary, the PadN option is used.)
- * Sub-Type is followed by 3 'x' bits, set to any value on transmission (typically all-zeros) and ignored on reception. Pad1 therefore consists of a single octet with the most significant 5 bits set to 0, and with no Sub-Length or Sub-Option Data fields following.

10.2.2. PadN

```

+-----+-----+-----+
| Sub-Type=1 | Sub-Length=N | N padding octets ...
+-----+-----+-----+

```

Figure 19: PadN

- * Sub-Type is set to 1.
- * Sub-Length is set to N that encodes the number of padding octets that follow.
- * Sub-Option Data consists of N octets, set to any value on transmission (typically all-zeros) and ignored on reception.

10.2.3. Node Identification

Nodes may include the Node Identification sub-option as supplementary identification information in addition to the IPv6 ND message Source Address. If multiple instances appear in the same OMNI option, the first instance of a specific ID-Type is processed and all other instances of the same ID-Type are ignored. (A single IPv6 ND message can therefore convey multiple distinct Node Identifications - each with a different ID-Type.)

The format and contents of the sub-option are shown in Figure 20:




Figure 20: Node Identification

- * Sub-Type is set to 2. Multiple instances are processed as discussed above.
- * Sub-Length is set to N that encodes the number of Sub-Option Data octets that follow. The ID-Type field is always present, and the maximum Node Identification Value length is limited by the remaining available space in this OMNI option.
- * ID-Type is a 2-octet field that encodes the type of the Node Identification Value. The following ID-Type values are currently defined:
 - 0 - Multilink Local Address (MLA). A special-purpose IPv6 address assigned to an OMNI interface for adaptation layer addressing as discussed in Section 8. Indicates that Node Identification Value contains a 16-octet MLA.
 - 1 - Universally Unique Identifier (UUID) [RFC9562]. Indicates that Node Identification Value contains a 16-octet UUID.
 - 2 - Network Access Identifier (NAI) [RFC7542]. Indicates that Node Identification Value contains an (N-1)-octet NAI.
 - 3 - Fully-Qualified Domain Name (FQDN) [RFC1035]. Indicates that Node Identification Value contains an (N-1)-octet FQDN.
 - 4 - IPv4 Address. Indicates that Node Identification contains a 4-octet IPv4 address. The IPv4 address type is determined with reference to the IANA IPv4 Address Space Registry [IPV4].
 - 5 - Unassigned.
 - 6 - IPv6 Address. Indicates that Node Identification contains a general-purpose 16-octet IPv6 address that is not an MLA. The IPv6 address type is determined according to the IPv6 addressing architecture [RFC4291] with reference to the IANA IPv6 Global Unicast Address Assignments Registry [IPV6].

- 7 - 65532 - Unassigned.
 - 65533 - 65534 - reserved for experimentation, as recommended in [RFC3692].
 - 65535 - reserved by IANA.
- * Node Identification Value is an (N-2)-octet field encoded according to the appropriate the "ID-Type" reference above.

OMNI interfaces code Node Identification Values used for DHCPv6 messaging purposes as a DHCP Unique Identifier (DUID) using the "DUID-EN for OMNI" format with enterprise number 45282 (see: Section 21) as shown in Figure 21:

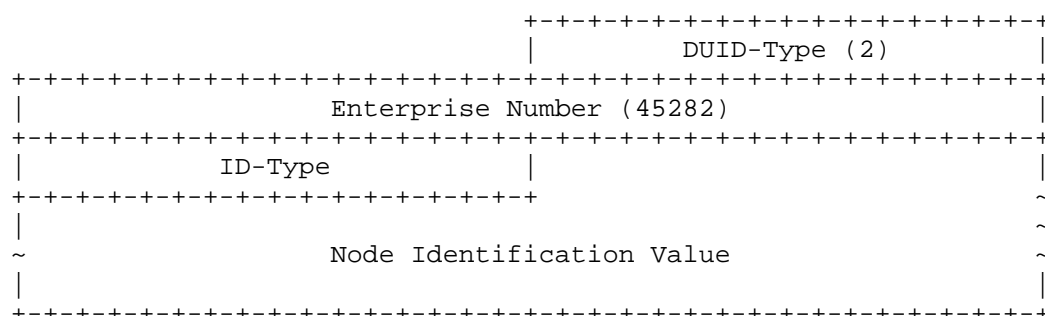


Figure 21: DUID-EN for OMNI Format

In this format, the OMNI interface codes the ID-Type and Node Identification Value fields from the OMNI sub-option following a 6-octet DUID-EN header, then includes the entire "DUID-EN for OMNI" in a DHCPv6 message per [I-D.ietf-dhc-rfc8415bis].

10.2.4. CGA

OMNI nodes include an OMNI Cryptographically Generated Address (CGA) sub-option for IPv6 ND messages the same as per Section 5.1 of [RFC3971] with the exception that the CGA body field itself need not be an integral number of 8-octet words. The OMNI CGA sub-option has the following format:




Figure 22: CGA

10.2.5. RSA Signature

The OMNI RSA Signature sub-option includes a public key-based signature extending over the length of the OMNI-encapsulated IPv6 ND message followed by any composite packet extensions then finally over the OMNI option itself up to but not including this sub-option. When present, the RSA Signature sub-option MUST appear as the first OMNI sub-option.

Each OMNI encapsulated IPv6 ND message should include at most one RSA Signature or HMAC sub-option. If an IPv6 ND message includes multiple RSA Signature and/or HMAC sub-options, the first one is processed and all others ignored.

The IPv6 ND message source calculates the IPv6 ND message checksum over the length of just the ND message itself and writes the value into the ICMPv6 Checksum field as a first step. The OAL source can then calculate a digital signature to include in an OMNI RSA Signature sub-option as discussed below. The OAL source finally calculates the OMNI option checksum and writes its value into the OMNI trailing Checksum1 field, then includes any trailing information and calculates/writes the Checksum2 field.

The RSA Signature sub-option is formatted as shown in Figure 23:




Figure 23: RSA Signature

- * Sub-Type is set to 4. The RSA Signature sub-option should appear at most once in any OMNI option; if multiple instances appear in the same OMNI option the final one is processed and all others are ignored.
- * Sub-Length is set to N, i.e., the length of the option in octets beginning immediately following the Sub-Length field and extending to the end of the Padding. The Key Hash is always 128 bits in length per [RFC3971] while the length of the Digital Signature is constrained by the remaining available space for this sub-option.
- * Key Hash, Digital Signature and Padding are included as specified in Section 5.2 of [RFC3971].

The sender constructs the Digital Signature in the same manner as Section 5.2 of [RFC3971] except over the following data:

1. For CGAs, the 128-bit CGA Message Type tag [RFC3972] value for SEND, 0x086F CA5E 10B2 00C9 9C8C E001 6427 7C08 [RFC3971]. For MLA types that include cryptographically generated elements, the 128-bit Context ID found in the "CGA Extension Type Tags" registry [IANA-CGA].
2. The entire IPv6 ND message beginning with the IPv6 header, then extending over the ND message header and all ND message options.
3. All composite IP packet extensions up to the beginning of the OMNI option.
4. All OMNI sub-options preceding this RSA Signature sub-option.

Note: the same as for ordinary IPv6 ND SEND operations, the CGA/MLA subject to authentication appears in the IPv6 ND message IPv6 Source Address. Note also that the OAL IPv6 Source and Destination Address as well as Payload Length are not included in the Digital Signature since these values may be rewritten by proxies on the path (i.e., both Proxy/Servers in different OMNI link segments and Proxy/Clients within the same OMNI link segment).

10.2.6. Timestamp

OMNI nodes include an OMNI Timestamp sub-option in IPv6 ND messages to ensure that unsolicited advertisements and redirects have not been replayed. The Timestamp sub-option is processed exactly the same as in Section 5.3.1 of [RFC3971]. The OMNI Timestamp sub-option has the following format:

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| Sub-Type=5 | Sub-length=8 |                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Reserved (6 octets)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Timestamp (8 octets)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 24: Timestamp

10.2.7. Nonce

OMNI nodes include an OMNI Nonce sub-option in IPv6 ND messages to ensure that an advertisement is a fresh response to a solicitation sent earlier by the node. The Nonce sub-option is processed exactly the same as in Section 5.3.2 of [RFC3971] with the exception that the Nonce field itself need not be an integral number of 8-octet words. The OMNI Nonce sub-option has the following format:

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| Sub-Type=6 | Sub-length=N |                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Nonce (N octets)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 25: Nonce

10.2.8. Trust Anchor

OMNI nodes include an OMNI Trust Anchor sub-option the same as described in Section 6.4.3 of [RFC3971] with the exception that the sub-option does not require 8-octet alignment and need not contain an integral number of 8 octet units. The Trust Anchor sub-option has the following format:

```

+-----+
| Sub-Type=7 | Sub-length=N | Name Type | Pad Length |
+-----+
|
~                Trust Anchor Body (N-2 octets)                ~
|
+-----+

```

Figure 26: Trust Anchor

10.2.9. Certificate

OMNI nodes include an OMNI Certificate sub-option in IPv6 ND messages the same as described in Section 6.4.4 of [RFC3971] with the exception that the sub-option does not require 8-octet alignment and need not contain an integral number of 8 octet units. The Certificate sub-option has the following format:

```

+-----+
| Sub-Type=8 | Sub-length=N | Cert Type | Reserved |
+-----+
|
~                Certificate Body (N-2 octets)                ~
|
+-----+

```

Figure 27: Certificate

10.2.10. Hashed Message Authentication Code (HMAC)

OMNI nodes may include a Hashed Message Authentication Code (HMAC) sub-option. When present, the HMAC sub-option must appear as the first sub-option the same as specified for RSA Signature above.

Each OMNI encapsulated IPv6 ND message should include at most one RSA Signature or HMAC sub-option. If an IPv6 ND message includes multiple RSA Signature and/or HMAC sub-options, the first one is processed and all others ignored.

The format of the HMAC option is taken directly from Section 2.1.2 of [RFC8754] as shown in Figure 28:

```

+-----+
| Sub-Type=9 | Sub-length=N |D|          RESERVED          |
+-----+
|          HMAC Key ID (4 octets)          |
+-----+
|          HMAC (variable)          |
~                                     ~
+-----+

```

Figure 28: Hashed Message Authentication Code (HMAC)

The HMAC sub-option is encoded and processed the same as specified in [RFC2104] and Section 2.1.2 of [RFC8754] except that the HMAC is applied over the following text:

1. For CGAs, the 128-bit CGA Message Type tag [RFC3972] value for SEND, 0x086F CA5E 10B2 00C9 9C8C E001 6427 7C08 [RFC3971]. For MLA types that include different cryptographically generated elements, the 128-bit Context ID found in the "CGA Extension Type Tags" registry [IANA-CGA].
2. The entire IPv6 ND message beginning with the IPv6 header, then extending over the ND message header and all ND message options.
3. All composite IP packet extensions up to the beginning of the OMNI option.
4. All OMNI sub-options following this HMAC sub-option.

Note: the same as for ordinary IPv6 ND SEND operations, the CGA/MLA subject to authentication appears in the IPv6 ND message IPv6 Source Address. Note also that the OAL IPv6 Source and Destination Address as well as Payload Length are not included in the Digital Signature since these values may be rewritten by proxies on the path (i.e., both Proxy/Servers in different OMNI link segments and Proxy/Clients within the same OMNI link segment).

10.2.11. Neighbor Synchronization

IPv6 ND messages that establish or update neighbor state between Clients and their Proxy/Servers or peer Clients include a Neighbor Synchronization OMNI sub-option. Each IPv6 ND message includes at most one Neighbor Synchronization sub-option which must be specific to the underlying interface pair over which ND messages are exchanged.

The Neighbor Synchronization sub-option is formatted as follows:

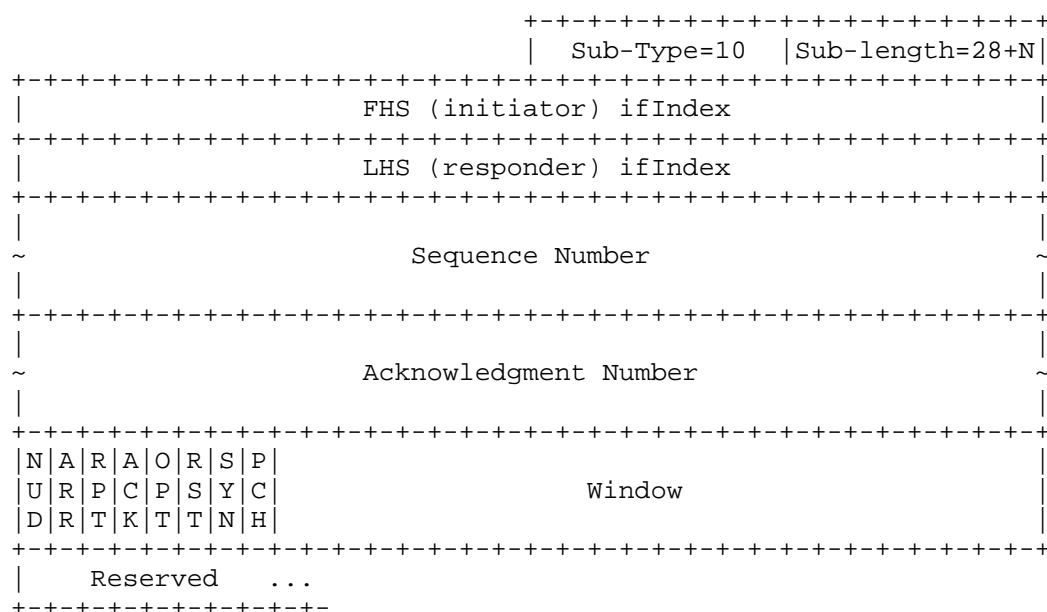


Figure 29: Neighbor Synchronization

- * Sub-Type is set to 10. If multiple instances appear in the same OMNI option, the first is processed and all others are ignored.
- * Sub-Length is set to $(28 + N)$, where N is the length in octets of the trailing Reserved field if present; otherwise, 0.
- * the first 8 octets include the 4-octet ifIndex of the FHS (initiator) node followed by the 4-octet ifIndex of the LHS (responder) node.
- * the next 20 octets of Sub-Option Data follows from the Transmission Control Protocol (TCP) header specified in Section 3.1 of [RFC9293]. The field is formatted as an 8-octet

Sequence Number, followed by an 8-octet Acknowledgement Number, followed by a 1-octet flags field followed by a 3-octet Window size. The TCP (ACK, RST, SYN) flags are used for TCP-like neighbor synchronization while intermediate nodes always cache the Sequence Number and Window size value regardless of the SYN flag setting as discussed in Section 6.7. The TCP (CWR, ECE, URG, PSH, FIN) flags are unused and replaced by the OMNI-specific flags (NUD, ARR, RPT, OPT, PCH).

- * Clients set the Neighbor Unreachability Detection (NUD), Address Resolution Responder (ARR) and Report (RPT) flags in RS messages to control the operation of their Proxy/Server neighbors as discussed in Section 13.
- * Neighbors set the OPT flag as discussed in Section 6.7 during a (SYN, SYN/ACK) synchronization exchange that does not require a concluding ACK.
- * OAL intermediate systems set the Path Change (PCH) flag in IPv6 ND control messages used to report a change in a path established by multilink forwarding.
- * An N-octet trailing Reserved field is available for expansion to include additional flags as necessary for future applications. Currently, no additional flags are defined and N should be set to 0.

10.2.12. Interface Attributes

The Interface Attributes sub-option provides neighbors with forwarding information for the multilink conceptual sending algorithm discussed in Section 12. Neighbors use the forwarding information to select among candidate underlay interfaces that can be used to forward carrier packets to the neighbor based on factors such as traffic selectors and link metrics. Interface Attributes further include link layer address information to be used for either direct INET encapsulation for targets in the local SRT segment or spanning tree forwarding for targets in remote SRT segments.

OMNI nodes include Interface Attributes for some/all of a source or target Client's underlay interfaces in IPv6 ND solicitation and response messages that exchange peer-to-peer Client information (see: [I-D.templin-6man-aero3]). At most one Interface Attributes sub-option for each distinct ifIndex may be included; if an IPv6 ND message includes multiple Interface Attributes sub-options for the same ifIndex, the first is processed and all others are ignored. OMNI nodes that receive IPv6 ND messages can use all of the included Interface Attributes and/or Traffic Selectors to formulate a map of the prospective source or target node as well as to seed the information to be populated in future neighbor exchanges.

OMNI Clients and Proxy/Servers also include Interface Attributes sub-options in RS/RA messages used to initialize, discover and populate routing and addressing information. Each RS message MUST contain exactly one Interface Attributes sub-option with an ifIndex corresponding to the Client's underlay interface used to transmit the message, and each RA message MUST echo the same Interface Attributes sub-option with any (proxied) information populated by the FHS Proxy/Server to provide operational context.

When an FHS Proxy/Server receives an RS message destined to an anycast L2 address, it MUST include an additional Interface Attributes sub-option with ifIndex '0' that encodes its own unicast L2 address relative to the Client's underlay interface in the solicited RA response. Any additional Interface Attributes sub-options that appear in RS/RA messages (i.e., besides those for the Client's own ifIndex and ifIndex '0') are ignored.

The Interface Attributes sub-option is formatted as shown below:

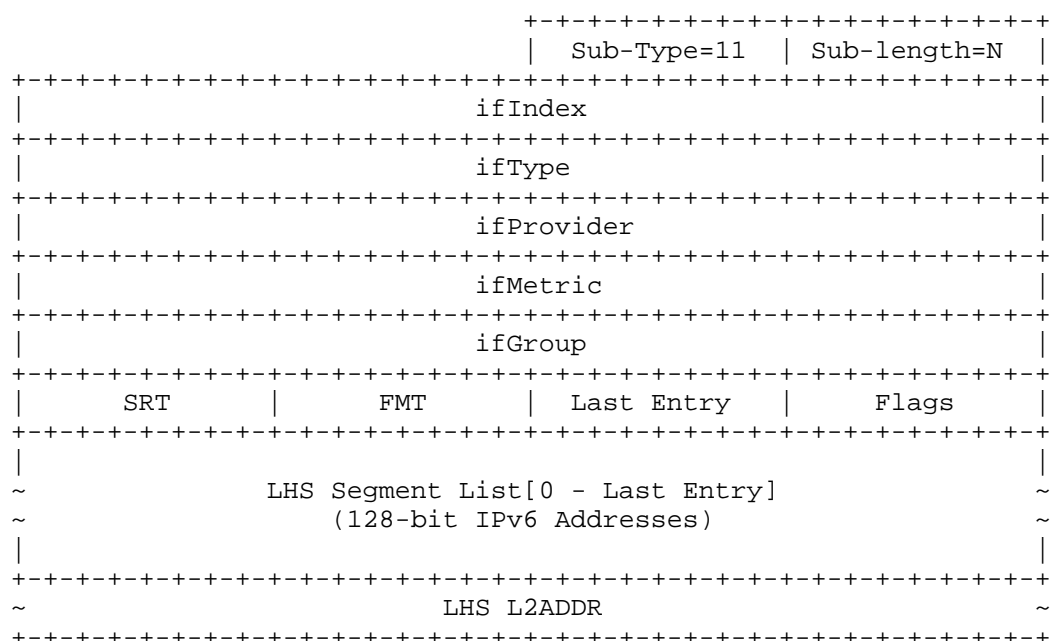


Figure 30: Interface Attributes

- * Sub-Type is set to 11. Multiple instances are processed as discussed above.
- * Sub-Length is set to N that encodes the number of Sub-Option Data octets that follow.
- * Sub-Option Data contains an "Interface Attributes" option encoded as follows:
 - ifIndex is a 4-octet index value corresponding to a specific underlay interface. Client OMNI interfaces MUST number each distinct underlay interface with a unique non-zero ifIndex value assigned by network management per [RFC2863] and include the value in this field. The ifIndex value '0' denotes "unspecified".
 - ifType is a 4-octet type value corresponding to this underlay interface. The value is coded per the 'IANAifType-MIB' registry [<http://www.iana.org>].

- ifProvider is a 4-octet provider identifier corresponding to this underlay interface. This document defines the single provider identifier value '0' (undefined). Future documents may define other values.
- ifMetric encodes a 4-octet interface metric. Lower values indicate higher priorities, and the highest value indicates an interface that should not be selected. The ifMetric setting provides an instantaneous indication of the interface bandwidth, link quality, signal strength, cost, etc.; hence, its value may change in successive IPv6 ND messages.
- ifGroup is a 4-octet identifier for a Link Aggregation Group (LAG) [IEEE802.1AX] corresponding to the underlay interface identified by ifIndex. Interface attributes for ifIndex members of the same group will encode the same value in ifGroup. This document defines the single ifGroup value '0' meaning "no group assigned". Future documents will specify the setting of other values.
- SRT is a 1-octet Segment Routing Topology prefix length that determines the length associated with this sub-tree of a larger Internetworking topology that may include the concatenation of multiple connected segments. Correspondent nodes apply the SRT prefix length to the Client's SNP GUA to discover a topological orientation for this interface.
- FMT - a 1-octet "Forward/Mode/Type" code interpreted as follows:
 - o The most significant 2 bits (i.e., "FMT-Forward" and "FMT-Mode") are interpreted in conjunction with one another. When FMT-Forward is clear, the LHS Proxy/Server performs OAL reassembly and decapsulation to obtain the original IP packet before forwarding. If the FMT-Mode bit is clear, the LHS Proxy/Server then forwards the original IP packet at L3; otherwise, it invokes the OAL to re-encapsulate, re-fragment and sends the resulting carrier packets to the Client via the selected underlay interface. When FMT-Forward is set, the LHS Proxy/Server forwards unmodified OAL fragments to the Client without reassembling. If FMT-Mode is clear, all carrier packets destined to the Client must always be sent via the LHS Proxy/Server; otherwise the Client is eligible for direct forwarding over the open INET where it may be located behind one or more NATs.

- o The least significant 6 bits ("FMT-Type") determines the type of L2 encapsulation needed to reach the target Client interface within its local *NET. When the most significant bit (msb) of FMT-Type is set, the interface has been determined to reside behind a Network Address Translator (NAT) as discovered during Client exchanges with their Proxy/Servers. The least significant 5 bits of FMT-Type encode an L2 encapsulation type value as follows:
 - + 0 - L2 encapsulation type is unspecified. No L2ADDR is included and the msb is ignored.
 - + 1 - Client interface is within a MANET where multihop forwarding occurs as an adaptation layer service. No L2ADDR is included and the msb is ignored.
 - + 2 - L2 encapsulation type is EUI-48 only. L2ADDR is 6 octets in length and encodes an EUI-48 address [EUI].
 - + 3 - L2 encapsulation type is EUI-64 only. L2ADDR is 8 octets in length and encodes an EUI-64 address [EUI].
 - + 4 - L2 encapsulation type is IPv4 only. L2ADDR is 4 octets in length and encodes an IPv4 address.
 - + 6 - L2 encapsulation type is IPv6 only. L2ADDR is 16 octets in length and encodes an IPv6 address.
 - + 7 - L2 encapsulation type is UDP/IPv4. L2ADDR is 6 octets in length and encodes a 4-octet IPv4 address followed by a 2-octet UDP port number.
 - + 8 - L2 encapsulation type is UDP/IPv6. L2ADDR is 18 octets in length and encodes a 16-octet IPv6 address followed by a 2-octet UDP port number.
 - + 5, [9 - 31] - Reserved for future use.
- Last Entry and Flags are 1-octet fields defined exactly as specified in [RFC8754]. In particular, Last Entry contains the index (zero based), in the LHS Segment List, of the last element of the list. Flags includes 8 Unused bits that must be 0 on transmission and ignored on receipt.
- LHS Segment List[0 - Last Entry] echoes the IPv6 addresses recorded by the LHS Client's prior RS message transmissions to the LHS Proxy/Server. The ordered list begins with the LHS Client's MLA as the first entry, followed by the MLAs of each

successive Proxy in the path to the Proxy/Server, and with the final entry including the SNP GUA delegated to the Client for this interface. The list provides context for the recipient to populate an SRH in subsequent IPv6 ND messages.

- LHS L2ADDR is N octets in length according to FMT-Type as discussed above. LHS L2ADDR identifies the LHS Client's *NET interface which may connect to an open INET or a private *NET behind one or more NATs. When L2ADDR includes an IPv4 or IPv6 address, it appears in network byte order in ones-complement "obfuscated" form per [RFC4380].
- The LHS information therefore satisfies per-interface address resolution and SRT/FMT/LHS together provide guidance for the OMNI interface forwarding algorithm. Specifically, if LHS::/SRT is located in the FHS OMNI link segment, then the source can address the target Client either through its dependent Proxy/Server or through direct L2 encapsulation (while engaging NAT traversal if necessary) according to FMT. Otherwise, the target Client is located on a different SRT segment and the path from the source must employ a combination of route optimization and spanning tree hop traversals.

10.2.13. Traffic Selector

The Traffic Selector sub-option provides forwarding information for the multilink conceptual sending algorithm discussed in Section 12. The sub-option includes an augmented traffic selector per [RFC6088] as ancillary information for an Interface Attributes sub-option with the same ifIndex value, or as discrete information for the included ifIndex when no Interface Attributes sub-option is present.

IPv6 ND messages may include multiple Traffic Selectors for some or all of the source/target Client's underlay interfaces (see: [I-D.templin-6man-aero3] for further discussion). Traffic Selectors must be honored by all implementations in the format shown below:




Figure 31: Traffic Selector

- * Sub-Type is set to 12. Multiple instances with the same or different ifIndex values may appear in the same OMNI option. When multiple instances appear, all are processed and the cumulative information from all is accepted.
- * Sub-Length is set to N that encodes the number of Sub-Option Data octets that follow.
- * Sub-Option data begins with a 4-octet ifIndex value corresponding to a specific underlay interface.
- * The remainder of Sub-Option Data contains one or more "Traffic Selector" blocks for this ifIndex that each begin with 1-octet "TS Format" and "TS Length" fields. TS length encodes the combined lengths of the Traffic Selector body that follows (i.e. up to 255

octets). When TS Format encodes the value 1 or 2, the Traffic Selector body encodes an IPv4 or IPv6 traffic selector per [RFC6088] beginning with 14 flag bits ("A-N"); when TS Format encodes any other value the Traffic Selector block is skipped and processing resumes beginning with the next Traffic Selector block (note that future specifications may define new TS Formats).

- * The Traffic Selector block elements then appear immediately after the flags (with no 16-bit Reserved field included) and encode the information corresponding to any set flag bit(s) in order the same as specified in [RFC6088]. Each included Traffic Selector block is processed consecutively, with its length subtracted from the remaining sub-option length until all blocks are processed. If the length of any Traffic Selector block would exceed the remaining length for the entire sub-option, the remainder of the sub-option is ignored.

10.2.14. Geo Coordinates




Figure 32: Geo Coordinates

- * Sub-Type is set to 13. If multiple instances appear in the same OMNI option all are processed.
- * Sub-Length is set to N that encodes the number of Sub-Option Data octets that follow.
- * Geo Type is a 1-octet field that encodes a type designator that determines the format and contents of the Geo Coordinates field that follows. The following types are currently defined:
 - 0 - NULL, i.e., the Geo Coordinates field is zero-length.
- * Geo Coordinates is a type-specific format field of length up to the remaining available space for this OMNI option. New formats to be specified in future documents and may include attributes such as latitude/longitude, altitude, heading, speed, etc.

10.2.15. Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Message

The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) sub-option may be included in the OMNI options of Client RS messages and Proxy/Server RA messages. The DHCPv6 sub-option is formatted per Section 8 of [I-D.ietf-dhc-rfc8415bis] as shown in Figure 33:




Figure 33: DHCPv6 Message

- * Sub-Type is set to 14. At most 2 instances may appear in a single OMNI option. If more instances appear, the first 2 are processed and all others are ignored.
- * Sub-Length is set to N that encodes the number of Sub-Option Data octets that follow. The 'msg-type' and 'transaction-id' fields are always present; hence, the length of the DHCPv6 options is limited by the remaining available space for this OMNI option.
- * 'msg-type' and 'transaction-id' are coded according to Section 8 of [I-D.ietf-dhc-rfc8415bis].
- * A set of DHCPv6 options coded according to Section 21 of [I-D.ietf-dhc-rfc8415bis] follows.

10.2.16. PIM-SM Message

The Protocol Independent Multicast - Sparse Mode (PIM-SM) Message sub-option may be included in the OMNI options of IPv6 ND messages. The PIM-SM message sub-option is formatted per Section 4.9 of [RFC7761] and as shown in Figure 34:




Figure 34: PIM-SM Message Option Format

- * Sub-Type is set to 15. If multiple instances appear in a single OMNI option all are processed.
- * Sub-Length is set to N, i.e., the length of the option in octets beginning immediately following the Sub-Length field and extending to the end of the PIM-SM message. The length of the entire PIM-SM message is therefore limited by the remaining available space for this OMNI option.
- * The PIM-SM message is coded exactly as specified in Section 4.9 of [RFC7761], except that the Checksum field is omitted since message integrity is already assured by the OMNI option checksum. The Reserved field is set to 0 on transmission and ignored on reception. The "PIM Ver" field encodes the value 2, and the "Type" field encodes the PIM message type. (See Section 4.9 of [RFC7761] for a list of PIM-SM message types and formats.)

10.2.17. Fragmentation Report (FRAGREP)

Fragmentation Report (FRAGREP) sub-options may be included in the OMNI options of unsolicited IPv6 ND control messages sent from an OAL destination to an OAL source on behalf of a specific flow. The message is formatted and processed the same as specified for the Fragmentation Report option in [I-D.templin-6man-ipid-ext2].

The message consists of the 20-bit Flow Label value for the source's flow, followed by the 11 most significant bits of the 16-bit Maximum Receive Unit (MRU) for this flow followed by a (L)oss indication. The MRU field is then followed by an Identification for the specific packet from the OAL source that triggered the flow plus an optional Bitmap marking the ordinal positions of individual fragments received and missing.

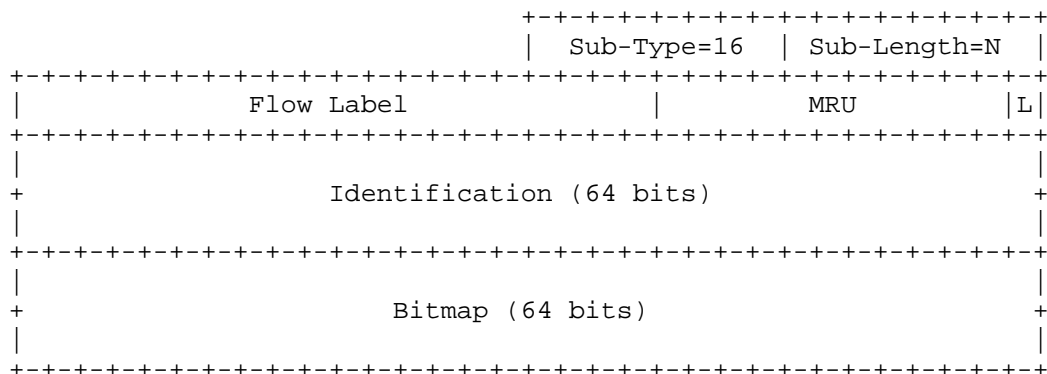


Figure 35: Fragmentation Report (FRAGREP)

- * Sub-Type is set to 16. If multiple instances appear in the same OMNI option all are processed.
- * Sub-Length is set to N which must be 20 if a Bitmap field is included or 12 otherwise.
- * Flow Label, MRU and L are 4 octets that include the same information as for the Fragmentation Report option in [I-D.templin-6man-ipid-ext2].
- * Identification includes the 8-octet Identification value found in a received OAL fragment.
- * Bitmap (optional) includes a 64-bit checklist of up to 64 ordinal fragments for this Identification, with each bit set to 1 for a fragment received or 0 for a fragment corrupted, lost or still in transit. For example, for a 20-fragment OAL packet with ordinal fragments #3, #10, #13 and #17 missing or corrupted and all other fragments received or still in transit, Bitmap(i) encodes the following:

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | 1 | 1 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 1 | 1 | 0 | 1 | 1 | 1 | 0 | 1 | 1 | 0 | 0 | 0 | ... |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 36

10.2.18. ICMPv6 Error

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| Sub-Type=17 | Sub-length=N | Type | Code |
+-----+-----+-----+-----+-----+-----+-----+
|                                     Type-Specific Data (4 octets)
+-----+-----+-----+-----+-----+-----+-----+
|                                     As much of invoking packet
+-----+-----+-----+-----+-----+-----+-----+
|                                     as possible without the IPv6 ND message
+-----+-----+-----+-----+-----+-----+-----+
|                                     exceeding the minimum IPv6 MTU [IPv6]
+-----+-----+-----+-----+-----+-----+-----+

```

Figure 37: ICMPv6 Error

- * Sub-Type is set to 17. If multiple instances appear in the same OMNI option all are processed.
- * Sub-Length is set to N that encodes the number of Sub-Option Data octets that follow.
- * Sub-Option Data includes an N-octet ICMPv6 Error Message body encoded per Section 2.1 of [RFC4443], with the ICMPv6 Checksum field omitted but with the full IPv6 header included in the invoking packet field, i.e., even if the message that elicited the error included a compressed header. (Note: ICMPv6 informational messages must not be included on transmission and must be ignored if received.)

10.2.19. Proxy/Server Departure

OMNI Clients include a Proxy/Server Departure sub-option in RS messages when they associate with a new FHS and/or MAP Proxy/Server and need to send a departure indication to an old FHS and/or MAP Proxy/Server. The Proxy/Server Departure sub-option is formatted as shown below:

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| Sub-Type=18 | Sub-length=32 |
+-----+-----+-----+-----+-----+-----+-----+
|                                     ~
|                                     Old FHS Proxy/Server L3ADDR (16 octets)
|                                     ~
+-----+-----+-----+-----+-----+-----+-----+
|                                     ~
|                                     Old MAP Proxy/Server L3ADDR (16 octets)
|                                     ~
+-----+-----+-----+-----+-----+-----+-----+

```

Figure 38: Proxy/Server Departure

- * Sub-Type is set to 18. If multiple instances appear in the same OMNI option, the first is processed and all others are ignored.
- * Sub-Length is set to 32.
- * Sub-Option Data contains the 16-octet L3ADDR for the "Old FHS Proxy/Server" followed by a 16-octet L3ADDR for an "Old MAP Proxy/Server". If the Old FHS/MAP is a different node, the corresponding L3ADDR includes the address of the (foreign) Proxy/Server. If the Old FHS/MAP is the local node, the corresponding L3ADDR includes the node's own address. If the FHS/MAP is unspecified, the corresponding L3ADDR instead includes the value ::/128.

11. Address Mapping - Multicast

The multicast address mapping of the native underlay interface applies. The Client mobile router also serves as an IGMP/MLD Proxy for its ENETs and/or hosted applications per [RFC4605].

The Client uses Multicast Listener Discovery (MLDv2) [RFC3810] to coordinate with Proxy/Servers, and underlay network elements use MLD snooping [RFC4541]. The Client can also employ multicast routing protocols to coordinate with network-based multicast sources as specified in [I-D.templin-6man-aero3].

Since the OMNI link model is NBMA, OMNI links support link-scoped multicast through iterative unicast transmissions to individual multicast group members (i.e., unicast/multicast emulation).

12. Multilink Conceptual Sending Algorithm

The Client's network layer selects the outbound OMNI interface according to SBM considerations when forwarding original IP packets from local or ENET applications to external correspondents. Each OMNI interface maintains an internal OAL neighbor cache maintained the same as discussed in [RFC4861], but also includes additional state for multilink coordination. Each Client OMNI interface maintains default routes via Proxy/Servers discovered as discussed in Section 13, and may configure more-specific routes discovered through means outside the scope of this specification.

For each original IP packet it forwards, the OMNI interface selects one or more source underlay interfaces based on PBM factors (e.g., traffic attributes, cost, performance, message size, etc.) and one or more target underlay interfaces for the neighbor based on Interface Attributes received in IPv6 ND messages (see: Section 10.2.11). Multilink forwarding may also direct carrier packet replication across multiple underlay interface pairs for increased reliability at

the expense of duplication. The set of all Interface Attributes and Traffic Selectors received in IPv6 ND messages determines the multilink forwarding profile for selecting target underlay interfaces.

When the OMNI interface forwards an original IP packet over a selected source underlay interface, it first employs OAL encapsulation and fragmentation as discussed in Section 5, then performs L2 encapsulation as directed by the appropriate AFV. The OMNI interface also performs L2 encapsulation (following OAL encapsulation) when the nearest Proxy/Server is located multiple hops away as discussed in Section 13.2.

OMNI interface multilink service designers MUST observe the BCP guidance in Section 15 [RFC3819] in terms of implications for reordering when original IP packets from the same flow may be spread across multiple underlay interfaces having diverse properties.

12.1. Multiple OMNI Interfaces

Clients may connect to multiple independent OMNI links within the same or different OMNI domains to support SBM. The Client configures a separate OMNI interface for each link so that multiple interfaces (e.g., omni0, omni1, omni2, etc.) are exposed to the network layer. Each OMNI interface is configured over a separate set of underlying interfaces and configures one or more OMNI link SRA addresses (see: Section 8); the Client injects the corresponding SRA prefixes into the ENET routing system. Multiple distinct OMNI links can therefore be used to support fault tolerance, load balancing, reliability, etc.

Applications in ENETs can use Segment Routing to select the desired OMNI interface based on SBM considerations. The application writes an OMNI link SRA address into the original IP packet's Destination Address, and writes the actual destination (along with any additional intermediate hops) into the Segment Routing Header. Standard IP routing directs the packet to the Client's mobile router entity, where the OMNI link SRA address identifies the correct OMNI interface for next hop forwarding. When the Client receives the packet, it replaces the IP Destination Address with the next Address found in the Segment Routing Header and forwards the message via the OMNI interface identified by the SRA address.

Note: The Client need not configure its OMNI interface indexes in one-to-one correspondence with the global OMNI Link-IDs configured for OMNI domain administration since the Client's indexes (i.e., omni0, omni1, omni2, etc.) are used only for its own local interface management.

12.2. Client-Proxy/Server Loop Prevention

After a Proxy/Server has registered an MNP for a Client (see: Section 13), the Proxy/Server will forward all original IP packets (or carrier packets) destined to an address within the MNP to the Client. The Client will under normal circumstances then forward the resulting original IP packet to the correct destination within its connected (downstream) ENETs.

If at some later time the Client loses state (e.g., after a reboot), it may begin returning original IP packets (or carrier packets) with destinations corresponding to its MNP to the Proxy/Server as its default router. The Proxy/Server therefore drops any original IP packets received from the Client with a Destination Address that corresponds to the Client's MNP (i.e., whether ULA or GUA), and drops any carrier packets with both Source and Destination Address corresponding to the same Client's MNP regardless of their origin.

Proxy/Servers support "hairpinning" for packets with SNP Source and Destination Addresses that would convey useful data from a source SNP Client to a target SNP Client both located in the same OMNI link segment. Proxy/Servers support this hairpinning according to [RFC6296], however ULA-to-ULA addressing between peer nodes within the same OMNI link segment is preferred whenever possible.

13. Router Discovery and Address/Prefix Delegation

Clients engage their FHS Proxy/Servers and the MS by sending OAL encapsulated RS messages with OMNI options under the assumption that one or more Proxy/Server will process the message and respond. The RS message is received by an FHS Proxy/Server, which may in turn forward a proxied copy to a MAP Proxy/Server located in a local or remote SRT segment if the Client requires MNP service. The MAP Proxy/Server then returns an OAL encapsulated RA message either directly to the Client or via the original FHS Proxy/Server acting as a proxy.

Clients send RS messages under the conditions specified in Section 6.3.7 of [RFC4861] which includes not only interface/link initialization conditions but also mobility factors. In particular, Clients may send RS messages when the OMNI interface or an underlay interface changes state, or when the Client moves to a new link and needs to discover addressing parameters for its new topological orientation. The Client's RS/RA exchanges therefore maintain the most current OMNI link state information even across frequent mobility events.

To initiate Router Discovery and Address/Prefix Delegation, the Client uses its OMNI interface LLA as the IPv6 Source Address for RS messages it sends to candidate FHS Proxy/Servers. The OMNI interface adaptation layer in turn translates the LLA into its MLA while also including OMNI authentication, Interface Attributes and any other sub-options. The FHS Proxy/Server first consults an address registration service to determine whether the Client is authorized to use its claimed MLA and discards the RS message if authorization fails. Otherwise, the FHS Proxy/Server provides Router Discovery and Address/Prefix Delegation services for the Client per the remainder of this section.

To support Client to service coordination, the OMNI option provides flag bits in the OMNI Neighbor Synchronization sub-option as discussed in Section 10.2.11. Clients set or clear Neighbor Synchronization flags in RS messages as directives to the Mobility Service FHS/MAP Proxy/Servers. Proxy/Servers interpret the flags as follows:

- * When an FHS Proxy/Server forwards or processes an RS with the NUD flag set, it responds directly to future NS Neighbor Unreachability Detection (NUD) messages with the Client as the target by returning NA(NUD) replies; otherwise, it forwards NS(NUD) messages to the Client.
- * When the MAP Proxy/Server receives an RS with the ARR flag set, it responds directly to future NS Address Resolution (AR) messages with the Client as the target by returning NA(AR) replies; otherwise, it forwards NS(AR) messages to the Client.
- * When the MAP Proxy/Server receives an RS with the RPT flag set, it maintains a Report List of recent NS(AR) message sources for the source or target Client and sends unsolicited IPv6 ND control messages to all list members if any aspects of the Client's underlay interfaces change.

Mobility Service Proxy/Servers function according to the NUD, ARR and RPT flag settings received in the most recent RS message to support dynamic Client updates.

Clients and FHS Proxy/Servers include an authentication signature as an OMNI sub-option in their RS/RA exchanges when necessary but always include valid IPv6 ND message and OMNI option checksums. Clients and Proxy/Servers use the information included in RS/RA messages to establish NCE state and OMNI link autoconfiguration information as discussed in this section.

For each underlay interface, the Client sends RS messages with OMNI options to coordinate with a (potentially) different FHS Proxy/Server for each interface but (normally) only with one MAP Proxy/Server at a given time. All Proxy/Servers accept original IP packets addressed to their LLAs or link-scoped multicast, OAL packets addressed to their anycast/unicast MLA/ULA/GUAs and carrier packets addressed to their anycast/unicast L2ADDRs. The Client typically selects one MAP Proxy/Server among any of its FHS Proxy/Servers, but may instead select any other Proxy/Server on the OMNI link.

Example L2ADDR discovery methods appear in [RFC5214] and include data link login parameters, name service lookups, static configuration, a DHCP option, a static "hosts" file, etc. In the absence of other information, the Client can resolve the DNS Fully-Qualified Domain Name (FQDN) "linkupnetworks.[domainname]" where "linkupnetworks" is a constant text string and "[domainname]" is a DNS suffix for the OMNI link (e.g., "example.com"). The name resolution will return a set of DNS resource records to populate a Potential Router List (PRL) that contains addresses of Proxy/Servers for the local OMNI link segment. When the underlay *NET does not support standard unicast server-based name resolution [RFC1035] the Client can engage a multicast service such as mDNS [RFC6762] within the local OMNI link segment.

Each FHS Proxy/Server configures an MLA and SNP ULA/GUA prefix pairs for the local OMNI link segment then advertises its L2ADDR(s) for discovery as above. Each Client can then manage its own SNP ULA/GUA addresses through DHCPv6 address autoconfiguration exchanges with FHS Proxy/Servers. The FHS Proxy/Servers discovered over multiple of the Client's underlay interfaces may configure the same or different SNP ULA/GUA prefix pairs, and the Client's ULA/GUA for each underlay interface will fall within the ULA/GUA for the OMNI link segment relative to each FHS Proxy/Server.

Clients configure OMNI interfaces that observe the properties discussed in previous sections. The OMNI interface and its underlay interfaces are said to be in either the "UP" or "DOWN" state according to administrative actions in conjunction with the interface connectivity status. An OMNI interface transitions to UP/DOWN through administrative action and/or through underlay interface state transitions. When a first underlay interface transitions to UP, the OMNI interface also transitions to UP. When all underlay interfaces transition to DOWN, the OMNI interface also transitions to DOWN.

When a Client OMNI interface transitions to UP, the IP layer sends an initial series of RS messages. The OMNI interface then appends a single OMNI option at the end of each RS message while sending an interface-specific copy of the message over each underlay interface. These OMNI RS messages will register an initial set of underlay

interfaces that are also UP and optionally also request an MNP delegation. The Client sends additional RS messages to refresh lifetimes and to register/deregister underlay interfaces as they transition to UP or DOWN. Guidelines for sending additional RS messages to generate corresponding RAs are found in Section 8.3.4 of [RFC5214], and are further extended to include proactive responses to mobility events.

The Client's OMNI interface sends initial RS messages over an UP underlay interface with source set to its LLA (otherwise with source set to the unspecified address (::/128) per [RFC4861]). The OMNI interface further sets the Destination Address to the LLA of a specific (MAP) Proxy/Server (otherwise to the link-scoped All-Routers multicast address ff02::2 [RFC4291]). The Client includes an OMNI option per Section 10 with a Neighbor Synchronization sub-option with the RS NUD, ARR and RPT flags set or cleared as necessary.

Clients also include an OMNI Neighbor Synchronization sub-option with FHS ifIndex set to the ifIndex of its own underlay interface and with LHS ifIndex set to 0 (i.e., the default ifIndex configured by all Proxy/Servers). The Client also sets Sequence Number to a randomly-chosen 8-octet value, sets AFVI to a randomly-chosen initial value and sets the Flow Label in the IPv6 header to 0. (If the Client needs to coordinate with a MAP Proxy/Server other than this FHS Proxy/Server, it also includes an SRH with the SNP SRA GUA for the MAP.) The resulting exchange will establish symmetric Identification windows for the Client and FHS Proxy/Server.

The Client next includes an Interface Attributes sub-option for the underlay interface with a NULL Segment List plus a DHCPv6 Solicit sub-option with IA_PD options. The Client then includes any other necessary OMNI sub-options such as Traffic Selectors, RSA Signature (or HMAC), Timestamp, Nonce, etc. The OMNI interface finally sets or clears the Interface Attributes FMT-Forward and FMT-Mode bits according to its desired FHS Proxy/Server service model as described in Section 10.2.11.

The Client next prepares to forward the RS over the underlay interface using OAL encapsulation. The Client's OMNI interface first changes the RS LLA Source Address to its own MLA and (if the RS Destination Address is an LLA unicast address) changes the RS Destination Address to the MLA of the FHS Proxy/Server. The OMNI interface then includes a certificate and authentication signature if necessary followed by the OMNI option trailing fields including the AFVI, Checksum1/2 and Null Segment List. The OMNI interface next optionally includes an SRH extension as specified above, sets the OAL Source Address to its own MLA and sets the OAL Destination Address to site-scoped All-Routers multicast (ff05::2) [RFC4291], the known FHS

Proxy/Server MLA or an anycast address. When L2 encapsulation is used, the Client next includes the discovered FHS Proxy/Server L2ADDR or an anycast address as the L2 destination then fragments if necessary and forwards the resulting carrier packet(s) into the underlay network. Note that the Client does not yet create a NCE, but instead caches its RS message transmissions in the OAL to match against any received RA messages.

When an FHS Proxy/Server receives a carrier packet containing an RS it sets aside the L2 and OAL headers then verifies the checksums/authentication signature while also consulting an MLA registration service based on the Client's claimed certificate. If the RS message authenticity/integrity is verified, the FHS Proxy/Server then creates/updates a NCE indexed by the Client's MLA. The FHS Proxy/Server then caches the OMNI Interface Attributes and any Traffic Selector sub-options while also caching the AFVI plus L2 (UDP/IP) and OAL Source/Destination Address information. The FHS Proxy/Server next examines the Client's MLA then coordinates with the local DHCPv6 server to either allocate a new SNP GUA or extend the lease lifetime for an existing Client SNP GUA.

The FHS Proxy/Server then generates an IPv6 address for the Client from its GUA prefix and proposes it in an IA_NA option of the DHCPv6 message for the local DHCPv6 server that includes the Client's MLA in a DUID option. A suitable method for address generation appears in [I-D.gont-dhcwg-dhcpv6-iids].) If the proposed address is unique (or already leased to this Client), the DHCPv6 Server will return success; otherwise, the FHS Proxy/Server generates new IPv6 addresses and repeats the DHCPv6 message exchange. The DHCPv6 address lease lifetime must be the same as the Router Lifetime reported in RA messages. The DHCPv6 lease lifetime must therefore be refreshed through additional RS/RA exchanges before Router Lifetime expires.

After successful DHCPv6 Address registration, the FHS Proxy/Server next caches the confirmed SNP ULA/GUAs in the (newly-created) NCE. The FHS Proxy/Server next caches the RS Neighbor Synchronization NUD flag and Neighbor Synchronization parameters if present (see: Section 10.1). If no SRH is included but an OMNI DHCPv6 sub-option with IA_PD options is present, the FHS Proxy/Server coordinates with the local DHCPv6 server for prefix delegation then assumes the MAP role as a default router entry point for injecting the Client's MNP(s) into the OMNI link routing system. The FHS/MAP Proxy/Server then caches the RS ARR and RPT flags to determine its role in processing NS(AR) messages and generating unsolicited IPv6 ND control messages (see: Section 10.1).

The FHS/MAP Proxy/Server then prepares to return an RA message directly to the Client by first populating the Cur Hop Limit, Flags, Router Lifetime, Reachable Time and Retrans Timer fields with values appropriate for the OMNI link. The FHS/MAP Proxy/Server next includes an OMNI option with a unique AFVI for this Client plus a Neighbor Synchronization sub-option with responsive window synchronization information. The FHS/MAP Proxy/Server also includes an authentication sub-option if necessary and a DHCPv6 Reply sub-option for the IA_PD option that was processed/populated by the DHCPv6 exchange. The Proxy Server then includes a copy of the Client's original Interface Attributes sub-option with its INET-facing interface information written in the FMT, SRT and LHS L2ADDR fields and also with the Segment List that was received in the RS message trailer. The Proxy/Server sets the final Segment List entry to the SNP GUA received in the DHCPv6 exchange and sets L2ADDR to the address it observes in the RS message L2 source address. If the Proxy/Server observes a different L2ADDR than the one supplied by the Client, it sets the NAT indication in FMT-Type.

The FHS/MAP Proxy/Server next sets or clears the FMT-Forward and FMT-Mode flags if necessary to convey its capabilities to the Client, noting that it should honor the Client's stated preferences for those parameters if possible or override otherwise. The FMT-Forward/Mode flags thereafter remain fixed unless and until a new RS/RA exchange establishes different values (see: Section 10.2.11 for further discussion). If the FHS/MAP Proxy/Server's Client-facing interface is different than its INET-facing interface, the Proxy/Server next includes a second Interface Attributes sub-option with ifIndex set to '0', with a unicast L2 address for its Client-facing interface in the L2ADDR field and with its SRA ULA in the Segment List.

The FHS/MAP Proxy/Server next then includes any other necessary OMNI sub-options such as Nonce, Timestamp, etc. The FHS/MAP Proxy/Server also includes any other necessary RA options including 2 PIOs to advertise the ULA/GUA SNP prefixes for the segment with (A=0; L=0) per [RFC8028]. The FHS/MAP Proxy/Server then sets the RA Source Address to its own LLA and sets the RA Destination Address to the RS Source Address. The FHS/MAP Proxy/Server's OMNI interface then changes the RA Source Address to its own MLA, calculates the authentication signature/checksums and performs OAL encapsulation while setting the OAL Source Address to its own MLA and Destination Address to the OAL Source Address that appeared in the RS. If the RS Segment List was non-null, the FHS/MAP Proxy/Server also includes an SRH that contains the MLAs of endpoint OAL intermediate systems on the reverse path from the Proxy/Server to the original Client (while resetting the OAL Destination Address accordingly). The FHS/MAP Proxy/Server then performs L2 encapsulation with L2 Source and Destination address information reversed from the RS L2 information and returns the resulting carrier packets to the Client over the same underlay interface the RS arrived on.

When an FHS Proxy/Server receives an RS with Destination Address set to link-scoped all-Routers multicast (ff02::2) or the MLA of a different Proxy/Server, with valid checksum/authentication signature and with an SRH supplied by the Client that contains an SNP SRA GUA, it acts as a proxy for a different Proxy/Server to serve as the MAP. The FHS Proxy/Server first locally processes the RS message the same as described above except that it does not process any DHCPv6 IA_PD options. The FHS Proxy/Server then sets the OAL Source Address to the Client's SNP GUA and sets the OAL Destination according to the RS message SRH provided by the Client. The FHS Proxy/Server next creates or updates a NCE for the Client (i.e., based on the Client's MLA) and caches the OAL Source Address, Neighbor Synchronization and Interface Attributes information as above.

The FHS Proxy/Server then clears the Neighbor Synchronization SYN/ACK/OPT flags and writes the RS L2ADDR and RS Segment List including the Client's SNP GUA into the corresponding Interface Attributes fields. The FHS Proxy/Server next calculates and includes the message checksums then performs L2 encapsulation and sends the resulting carrier packets into the SRT secured spanning tree.

When the MAP Proxy/Server receives the carrier packet, it performs L2 decapsulation and OAL decapsulation to obtain the proxied RS, verifies the checksums, then performs DHCPv6 IA_PD processing to obtain or update any MNPs for the Client. The MAP Proxy/Server then creates/updates a NCE indexed by the Client's MLA and caches any state (including delegated MNPs, the ARR and RPT flags, IA_NA addresses, OAL addresses, Interface Attributes information and Traffic Selectors), then finally injects any delegated MNPs into the OMNI link routing protocol.

The MAP Proxy/Server then returns an OMNI encapsulated RA that echoes the Client's (proxied) Interface Attributes sub-option and with any RA parameters the same as specified for the FHS/MAP Proxy/Server case above while also including a DHCPv6 Reply sub-option with the IA_PD transaction results. The MAP Proxy/Server sets the RA Source Address to its own MLA and sets the Destination Address to the RS Source Address. The OMNI interface of the MAP Proxy/Server then sets the OAL Source Address to its own SNP SRA GUA and Destination Address to the cached value for the RS source. The MAP Proxy/Server then calculates the message checksums and encapsulates the RA as an OAL packet. The MAP Proxy/Server finally performs L2 encapsulation and sends the resulting carrier packet into the secured spanning tree.

When the FHS Proxy/Server receives the carrier packet it performs L2 decapsulation followed by OAL decapsulation to obtain the RA message, verifies checksums then updates the OMNI interface NCE for the Client and creates/updates a NCE for the MAP. The FHS Proxy/Server then sets the P flag in the RA flags field [RFC4389] and proxys the RA by changing the OAL source to its MLA and changing the OAL Destination Address to the Source Address from the Client's original RS message while also recording the Client's SNP ULA/GUA address pair as alternate indexes into the Client NCE. The FHS Proxy/Server then includes 2 RA message PIOs with (A=0; L=0) with the SNP ULA/GUA prefixes for the segment per [RFC8028].

The FHS Proxy/Server next includes a Neighbor Synchronization sub-option with its responses to its cached initiations from the Client. The FHS Proxy/Server also includes an Interface Attributes sub-option with ifIndex '0' and with its Client-facing interface unicast L2 address if necessary (see above) and includes an RSA Signature or HMAC sub-option if necessary. The FHS Proxy/Server next includes a unique AFVI for this Client then calculates the authentication signature and message checksums. The FHS Proxy/Server then includes an SRH extension to the OAL header if necessary with the MLAs of endpoint intermediate systems on the reverse path to the Client and with the OAL Destination Address adjusted accordingly. The FHS Proxy/Server finally performs L2 encapsulation with L2 addresses taken from the Client's NCE and sends the resulting carrier packet via the same underlay interface over which the RS was received.

When the Client receives the carrier packet, it performs L2 decapsulation followed by OAL decapsulation to obtain the RA message. The Client next verifies the authentication signature/checksums, then matches the RA with its previously-sent RS by comparing the RS Sequence Number with the RA Acknowledgement Number and also comparing the Nonce and/or Timestamp values. If the values match, the Client then creates/updates OMNI interface NCEs for both the MAP and FHS Proxy/Server and caches the information in the RA message. The Client next discovers its own SNP GUA address by examining the proxied Interface Attributes sub-option and discovers the SNP ULA/GUA PIO prefixes for the OMNI link segment per [RFC8028].

The Client then adds the ULA/GUA prefixes to the OMNI interface Prefix List associated with this FHS Proxy/Server and regards the corresponding ULA/GUA SRA addresses as the Proxy/Server addresses. If the Client has multiple underlay interfaces, it creates additional FHS Proxy/Server NCEs as necessary when it receives RAs over those interfaces (noting that multiple of the Client's underlay interfaces may be serviced by the same or different FHS Proxy/Servers). If the RA message includes a DHCPv6 Reply with the results of an IA_PD transaction, the Client provisions the delegated prefixes on downstream-facing links. The network layer of the Client finally adds each FHS Proxy/Server LLA (i.e., as determined by the adaptation layer mapping of the MLA) to the OMNI interface Default Router List.

For each underlay interface, the Client next caches the (filled-out) Interface Attributes for its own ifIndex including the L2ADDR and Segment List information that it received in an RA message over that interface so that it can include them in future IPv6 ND messages to provide neighbors with accurate FMT/SRT/LHS information. (If the message includes an Interface Attributes sub-option with ifIndex '0', the Client also caches the L2ADDR as the underlay network-local unicast address of the FHS Proxy/Server via that underlay interface.) The Client then consults the FMT-Type and L2ADDR to determine whether there may be NATs on the path to the FHS Proxy/Server.

The Client then caches the Neighbor Synchronization responsive window synchronization parameters for use in future IPv6 ND message exchanges via this FHS Proxy/Server. The Client finally configures default routes and assigns the IPv6 SRA address corresponding to the MNP (e.g., 2001:db8:1:2::) to a downstream network interface. The Client's OMNI interface then forwards the RA message to the IP layer which can then update its view of the neighbor cache and default router list.

Following the initial exchange, the FHS Proxy/Server MAY later send additional periodic and/or event-driven unsolicited RA messages per [RFC4861]. (The unsolicited RAs may be initiated either by the FHS Proxy/Server itself or by the MAP via the FHS as a proxy.) The Client then continuously manages its underlay interfaces according to their states as follows:

- * When an underlay interface transitions to UP, the Client sends an RS over the underlay interface with an OMNI option with sub-options as specified above.
- * When an underlay interface transitions to DOWN, the Client sends unsolicited IPv6 ND control messages over any UP underlay interface with an OMNI option containing Interface Attributes sub-options for the DOWN underlay interface with ifMetric set to 'ffffffff'. The Client sends isolated unsolicited IPv6 ND control messages when reliability is not thought to be a concern (e.g., if redundant transmissions are sent on multiple underlay interfaces), or may instead send an IPv6 ND solicitation message to receive a solicited reply.
- * When the Router Lifetime for the MAP Proxy/Server nears expiration, the Client sends an RS over any underlay interface to receive a fresh RA from the MAP. If no RA messages are received over a first underlay interface (i.e., after retrying), the Client marks the underlay interface as DOWN and should attempt to contact the MAP Proxy/Server via a different underlay interface. If the MAP Proxy/Server is unresponsive over additional underlay

interfaces, the Client sends an RS message with Destination Address set to the MLA of another Proxy/Server which will then assume the MAP role.

- * When all of a Client's underlay interfaces have transitioned to DOWN (or if a prefix delegation lifetime expires), the MAP Proxy/Server withdraws the MNP the same as if it had received a message with a release indication.

The Client is responsible for retrying each RS exchange up to MAX_RTR_SOLICITATIONS times separated by RTR_SOLICITATION_INTERVAL seconds until an RA is received. If no RA is received over an UP underlay interface (i.e., even after attempting to contact alternate Proxy/Servers), the Client can either declare this underlay interface as DOWN or continue to use the interface to support any peer-to-peer local communications with peers located in the same *NET. When changing to a new FHS/MAP Proxy/Server, the Client also includes a Proxy/Server Departure OMNI sub-option in new RS messages; the (new) FHS Proxy/Server will in turn send unsolicited IPv6 ND control messages to the old FHS and/or MAP Proxy/Server to announce the Client's departure as discussed in [I-D.templin-6man-aero3].

The network layer engages the OMNI interface as an ordinary IPv6 interface. Therefore, when the network layer sends an RS message the OMNI interface eventually returns corresponding RA messages from each responding FHS Proxy/Server. Each RA message contains configuration information consistent with the information received from the RAs generated by the Proxy/Servers. Note that this same logic applies to IPv4 implementations that employ "ICMP Router Discovery" per [RFC1256]; the OAL must convert ICMPv4 RS/RA messages into IPv6 ND RS/RA messages in a manner outside the scope of this specification prior to OMNI encapsulation.

Note: The Router Lifetime value in RA messages indicates the time before which the Client must send another RS message over this underlay interface (e.g., 600 seconds), however that timescale may be significantly longer than the lifetime the MS has committed to retain the prefix registration (e.g., REACHABLE_TIME seconds). Proxy/Servers are therefore responsible for updating MS state on a shorter timescale than the Client may be required to do on its own behalf.

Note: On certain multicast-capable underlay interfaces, Clients should send periodic unsolicited multicast NA messages and Proxy/Servers should send periodic unsolicited multicast RA messages as "beacons" that can be heard by other nodes on the link. If a node fails to receive a beacon after a timeout value specific to the link, it can initiate Neighbor Unreachability Detection (NUD) exchanges to test reachability.

Note: Although the Client's FHS Proxy/Server is a first-hop segment node from its own perspective, the Client stores the Proxy/Server's FMT, SRT, and addresses as last-hop segment (LHS) information to supply to neighbors. This allows both the Client and MAP Proxy/Server to supply the information to neighbors that will perceive it as LHS information on the return path to the Client.

Note: The MAP Proxy/Server injects Client MNPs into the OMNI link routing system by simply creating a route-to-interface forwarding table entry for MNP::/N via the OMNI interface. The dynamic routing protocol will notice the new entry and propagate the route to its peers. If the MAP receives additional RS messages, it need not re-create the forwarding table entry (nor disturb the dynamic routing protocol) if an entry is already present. If the MAP ceases to receive RS messages from any of the Client's interfaces, it removes the Client MNP(s) from the forwarding table (i.e., after a short delay) which also results in their removal from the routing system.

Note: If the Client's initial RS message includes an anycast L2 Destination Address, the FHS Proxy/Server returns the solicited RA using the same anycast address as the L2 source while including an Interface Attributes sub-option with ifIndex '0' and its true unicast address in the L2ADDR. When the Client sends additional RS messages, it includes this FHS Proxy/Server unicast address as the L2 Destination Address and the FHS Proxy/Server returns the solicited RA using the same unicast address as the L2 Source Address. This will ensure that RS/RA exchanges are not impeded by any NATs on the path while avoiding long-term exposure of messages that use an anycast address as the source.

Note: Clients should set the NUD, ARR and RPT flags consistently in successive RS messages and only change those settings when an FHS/MAP Proxy/Server service profile update is necessary.

13.1.1. Client-Proxy/Server Window Synchronization

The RS/RA exchanges discussed above observe the principles specified in Section 6.7. Window synchronization is conducted between the Client and each FHS Proxy/Server used to contact the MAP Proxy/Server, i.e., and not between the Client and the MAP. This is due to the fact that the MAP Proxy/Server is responsible only for forwarding messages via the secured spanning tree to FHS Proxy/Servers, and is not responsible for forwarding messages directly to the Client over unsecured networks.

When a Client sends an RS to perform window synchronization via a new FHS Proxy/Server, it includes an OMNI Neighbor Synchronization sub-option with window synchronization parameters with FHS ifIndex set to

its own interface index, with LHS ifIndex set to 0, with the SYN flag set and ACK flag clear, and with an initial Sequence Number. The Client also includes an Interface Attributes sub-option then performs OAL encapsulation and L2 encapsulation and sends the resulting carrier packet to the FHS Proxy/Server. When the FHS Proxy/Server receives the carrier packet, it performs L2 decapsulation, then extracts the RS message and caches the Neighbor Synchronization parameters. In the process, the FHS Proxy/Server removes the Neighbor Synchronization sub-option itself, since the path to the MAP Proxy/Server is not included in window synchronization.

The FHS Proxy/Server then performs L2 encapsulation and sends the resulting carrier packet via the secured spanning tree to the MAP Proxy/Server, which updates the Client's Interface Attributes and returns a unicast RA message. The MAP Proxy/Server performs OAL encapsulation followed by L2 encapsulation and sends the resulting carrier packet via the secured spanning tree to the FHS Proxy/Server. The FHS Proxy/Server then proxys the message as discussed in the previous section and includes a Neighbor Synchronization sub-option with responsive window synchronization information. The FHS Proxy/Server then forwards the message to the Client via OAL encapsulation which updates its window synchronization information for the FHS Proxy/Server as necessary.

Following the initial RS/RA-driven window synchronization, the Client can re-assert new windows with specific FHS Proxy/Servers by performing RS/RA exchanges between its own MLA and the MLAs of the FHS Proxy/Servers at any time without having to disturb the MAP. When the Client also needs to refresh MAP state, it can set the RS Destination Address to the MAP MLA and include an SRH if necessary to support FHS Proxy/Server RS forwarding.

This window synchronization is necessary only for MANET and INET Clients that must include authentication signatures with their IPv6 ND messages; Clients in secured ANETs can omit window synchronization. When Client-to-Proxy/Server window synchronization is used, subsequent IPv6 ND messages exchanged between peers include IPv6 Extended Fragment Headers in the OAL encapsulations with in-window Identification values to support message authentication. No header compression state is maintained by OAL intermediate systems, which only maintain state for per-flow data plane windows.

13.2. Router Discovery in IP Multihop and IPv4-Only Networks

On some *NETs, a Client may be located multiple intermediate OAL hops away from the nearest OMNI link Proxy/Server. Clients in multihop networks perform route discovery through the application of an adaptation layer routing protocol (e.g., a MANET routing protocol over omnidirectional wireless interfaces, etc.). Example routing protocols optimized for MANET operations include OSPFv3 [RFC5340] with MANET Designated Router (OSPF-MDR) extensions [RFC5614], OLSRv2 [RFC7181], AODVv2 [I-D.perkins-manet-aodvv2] and others. Clients employ the routing protocol according to the link model found in [RFC5889] and subnet model articulated in [RFC5942]. For unique identification within the MANET, Clients use an MLA as a Router ID.

MANETs can be compartmentalized internally with some nodes configured as simple Clients and others (that may have both "upstream" and "downstream" underlay interfaces) configured as cluster heads that act as Proxy/Clients. The Proxy/Clients configure and listen to the same multicast and anycast addresses as for Proxy/Servers on their downstream interfaces in order to act as endpoint OAL intermediate node proxys for other downstream Clients. Clusters within clusters based on these cluster head Proxy/Clients can then be recursively nested to multiple depths as long as at least one ultimate Proxy/Client configures an upstream interface that can directly address a Proxy/Server with connectivity to the outside Internetwork.

A Client located potentially multiple OAL hops away from the nearest Proxy/Server prepares an RS message, sets the Source Address to its MLA, and sets the Destination Address to link-scoped All-Routers multicast (ff02::2) or the MLA of a Proxy/Client or Proxy/Server as discussed above. The OMNI interface then employs OAL encapsulation, sets the OAL Source Address to its MLA and sets the OAL Destination Address to the MLA of the Proxy, the site-scoped All-Routers multicast address (ff05::2) or the OMNI IPv6 anycast address.

For IPv6-enabled *NETs where the underlay interface observes the MANET properties discussed above, the Client injects the MLA into the IPv6 multihop routing system and forwards the RS message without further encapsulation. Otherwise, the Client encapsulates the message in UDP/IPv6 L2 headers, sets the L2 Source Address to the underlay interface IPv6 address and sets the L2 Destination Address to the discovered unicast or anycast address of a Proxy. The Client then forwards the message into the IPv6 multihop routing system which conveys it to the nearest Proxy.

For IPv4-only *NETs, the Client encapsulates the RS message in UDP/IPv4 L2 headers, sets the L2 Source Address to the underlay interface IPv4 address and sets the L2 Destination Address to the discovered

unicast address of a Proxy/Server or the OMNI IPv4 anycast address. The Client then forwards the message into the IPv4 multihop routing system which conveys it to the nearest Proxy/Server advertising the corresponding IPv4 prefix. If the nearest Proxy/Server is too busy, it should forward (without Proxying) the OAL-encapsulated RS to another nearby Proxy/Server connected to the same IPv4 (multihop) network that configures the OMNI IPv6 anycast address. (In environments where reciprocal RS forwarding cannot be supported, the first Proxy/Server should instead return an RA based on its own MSP(s).)

When an OAL intermediate node that participates in the routing protocol receives the encapsulated RS, it forwards the message according to its OAL IPv6 forwarding table (note that an OAL intermediate system could be a fixed infrastructure element such as a roadside unit or another MANET/VANET Client). This process repeats iteratively until the RS message is received by a penultimate OAL hop within single-hop communications range of a Proxy/Server, which forwards the message to the Proxy/Server final hop.

When a Proxy/Server that configures the OMNI IPv6 anycast address receives the message, it decapsulates the RS and assumes either the MAP or FHS role (in which case, it may forward the RS to a candidate MAP). The MAP/FHS Proxy/Server then prepares an RA message using the same addressing disciplines as discussed in Section 13 and forwards the RA either to the FHS Proxy/Server or directly to the Client.

When the MAP or FHS Proxy/Server forwards the RA to the Client, it encapsulates the message in L2 encapsulation headers if necessary. The Proxy/Server then forwards the message to an OAL node within communications range, which forwards the message according to the next OAL hop by consulting its OAL IPv6 forwarding tables. The multihop forwarding process within the *NET continues repetitively until the message arrives at the original Client, which decapsulates the message and performs autoconfiguration the same as if it had received the RA directly from a Proxy/Server on the same physical link. The Client can optionally inject the delegated ULA/GUA and any MNP SRA GUAs into the IPv6 multihop routing system but this may cause excessive routing protocol overhead in some networks.

MANETs often include Clients that configure multiple interfaces, with downstream interfaces internal to the MANET and upstream interfaces connected to external *NETs. Such Clients can provide proxy services to enable router discovery for peer Clients that connect only internally within the MANET. These Proxy/Clients first perform router discovery to associate with true Proxy/Servers located on upstream *NETs. The Proxy/Clients also subscribe to the site-scoped all-routers multicast group (i.e., ff05::2) and advertise reachability for the OMNI IPv6 anycast address over their MANET interfaces.

When a source Client sends an initial RS message seeking service, MANET routing will direct the RS to one or more nearby Proxy/Clients which in turn forward the RS to one or more upstream interface Proxys. Each such Proxy/Client writes its MLA as the final Segment List IPv6 address at the end of the RS OMNI option trailer. The natural progression continues from innermost Proxy/Clients to outermost Proxy/Clients until the RS message reaches a Proxy/Server. By that time, the Segment List at the end of the RS OMNI option trailer will contain an ordered list of MLAs of all Proxy/Clients in the reverse path.

The MANET Proxy/Client model recursively extends to include arbitrarily many layers of nested MANETs between the source Client and external Proxy/Servers. When the source Client's first-hop Proxy/Client forwards an RS, it updates an adaptation layer neighbor cache entry for this Client's RS Source Address to include the OAL address of both the previous hop downstream (Proxy/)Client and next hop upstream (Proxy/)Client. The Proxy/Client then changes the OAL Source Address to its own MLA and forwards the RS to the next upstream Proxy/Client in succession which also updates state and changes the OAL Source Address to its own MLA. The progression continues until the RS reaches an ultimate upstream Proxy/Client that can directly contact a Proxy/Server via L2 encapsulation over an upstream *NET interface.

The Proxy/Server processes the RS and returns an RA while including its own MLA in the OAL Source Address and the MLA of the outermost Proxy/Client in the OAL Destination Address. The Proxy/Server also includes an SRH with the ordered list of Proxy/Client MLAs received from the RS message plus the MLA of the original source Client as the ultimate Segment List entry. When the outermost Proxy/Client receives the RA, it forwards the message to the MLA of the next hop Proxy/Client in succession based on the SRH information until the message arrives at the source Client. The source Client can then update its SNP GUA/ULA addresses, prefix list and default router list based on the information returned by the Proxy/Server. The Client also retains the Interface Attributes Segment List for future peer address resolution and multihop forwarding purposes.

When the Proxy/Server returns an RA, each upstream Proxy/Client forwards the RA through the recursively descending chain of downstream Proxy/Clients on the path to the source Client. Each Proxy/Client rewrites the OAL Destination Address according to the SRH next hop MLA address for the next downstream Proxy/Client hop toward the source Client and rewrites the OAL Source Address to its own MLA while caching the OAL Source Address from the previous upstream hop. When the RA arrives at the source Client, all upstream Proxy/Clients in the path will have established the state necessary for future packet forwarding.

Note that this service model applies equally for MANETs that have only Proxy/Client access to external *NET Proxy/Servers as well as those that have some mix of Proxy/Clients and true Proxy/Servers at the MANET border. True Proxy/Servers at the MANET border will service MANET Client router discovery requests the same as for any *NET, while external Proxy/Servers will discover potentially many MANET Clients all using the same L2ADDR belonging to a single Proxy/Client. This arrangement ensures that MANET-internal Clients are able to access external Internetworking services the same as for MANET border Clients that also have direct connections to external *NETs.

Note: When the RS message includes an anycast L2 encapsulation Destination Address, the FHS Proxy/Server must use the same anycast addresses as the L2 encapsulation Source Address to support forwarding of the RA message plus any initial data messages. The FHS Proxy/Server then sends the resulting carrier packets over any NATs on the path. When the outermost (Proxy/)Client receives the RA, it will discover the FHS Proxy/Server unicast L2 encapsulation address and can send future carrier packets using the unicast (instead of anycast) addresses to populate NAT state in the forward path. (If the Client does not have immediate data to send to the FHS Proxy/Server, it can instead send an OAL "bubble" - see Section 6.11.)

After the Client begins using unicast L2 encapsulation addresses in this way, the FHS Proxy/Server should also begin using the same unicast address in the reverse direction.

Note: When an OMNI interface configures an MLA, any nodes that forward an encapsulated RS message with the MLA as the OAL source must not consider the message as being specific to a particular OMNI link segment. MLAs can therefore also serve as the Source and Destination Addresses of unencapsulated IPv6 data communications within the local routing region; if the MLAs are injected into the local network routing protocol their prefix length must be set to 128 per [RFC5889].

13.3. DHCPv6-based Prefix Registration

When a Client requires SNP ULA/GUA delegations via a specific Proxy/Server (or, when the Client requires MNP delegations for the OMNI link), it invokes the DHCPv6 service [I-D.ietf-dhc-rfc8415bis] in conjunction with its OMNI RS/RA message exchanges.

When a Client requires the MS to delegate PA ULA/GUA pairs or PI MNPs, it sends an RS message to an FHS Proxy/Server that includes an Interface Attributes OMNI sub-option with a Segment List with final entry containing a previously-delegated SNP GUA or `::/128`. If the Client also requires one or more MNP delegations, it includes an OMNI DHCPv6 Message sub-option for MNP delegations. Each DHCPv6 sub-option contains a Client Identifier, an IA_PD option and a Rapid Commit option then sets the 'msg-type' field to "Solicit" and includes a 3-octet 'transaction-id'. The Client then sets the RS Destination Address to link-scoped All-Routers multicast (`ff02::2`) and sends the message using OAL encapsulation as discussed above.

When the FHS/MAP Proxy/Server receives the RS message, it examines the OMNI option contents. Next, if the OMNI option includes an Interface Attributes sub-option the FHS/MAP Proxy/Server acts as a "Proxy DHCPv6 Client" in a self-generated DHCPv6 IA_NA message exchange with the locally-resident DHCPv6 server while supplying the MLA of the Client as a DUID. When the Interface Attributes SNP GUA is unspecified (`::/128`), the FHS/MAP should first generate an SNP GUA that is likely to be unique using a suitable method such as that proposed in [I-D.gont-dhcwg-dhcpv6-iids]. The FHS/MAP Proxy/Server then includes the SNP GUA in an IA_NA option and sends the DHCPv6 message to the DHCPv6 Server, which verifies the SNP GUA and returns a DHCPv6 Reply message with autoconfiguration parameters. Next, if the Proxy/Server will act as a MAP it processes any DHCPv6 sub-options with an IA_PD message and performs a 2-message DHCPv6 PD exchange with the local DHCPv6 server exactly as specified in [I-D.ietf-dhc-rfc8415bis].

When the FHS Proxy/Server receives a DHCPv6 Reply with delegated addresses, it records the delegated SNP GUA plus a ULA from the companion ULA prefix that uses the same interface identifier in the NCE for the Client, then forwards the RS message to the MAP Proxy/Server for prefix delegation if necessary; otherwise, it returns an immediate RA message to the Client.

When the MAP Proxy/Server receives a DHCPv6 Reply with delegated prefixes, it creates OMNI interface MNP forwarding table entries (i.e., to prompt the dynamic routing protocol). The MAP Proxy/Server then sends an RA back to the FHS Proxy/Server with the Client's filled-out Interface Attributes sub-option and the DHCPv6 Reply message for the IA_PD delegation included in an OMNI DHCPv6 message sub-option. The FHS Proxy/Server then returns the RA to the Client.

14. Secure Redirection

If the *NET link model is multiple access, the FHS Proxy/Server is responsible for assuring that address duplication cannot corrupt the neighbor caches of other nodes on the link through the use of the DHCPv6 address delegation service. When the Client sends an RS message on a multiple access *NET, the Proxy/Server verifies that the Client is authorized to use the address and responds with an RA (or forwards the RS to the MAP) only if the Client is authorized.

After verifying Client authorization and returning an RA, the Proxy/Server MAY return IPv6 ND Redirect messages in response to subsequent data plane packet transmissions to direct Clients located on the same *NET to exchange OAL packets directly without transiting the Proxy/Server. In that case, the Clients can exchange OAL packets according to their unicast L2 addresses discovered from the Redirect message instead of using the dogleg path through the Proxy/Server. In some *NETs, however, such direct communications may be undesirable and continued use of the dogleg path through the Proxy/Server may provide better performance. In that case, the Proxy/Server can refrain from sending Redirects, and/or Clients can ignore them.

15. Proxy/Server Resilience

*NETs SHOULD deploy Proxy/Servers in Virtual Router Redundancy Protocol (VRRP) [RFC5798] configurations so that service continuity is maintained even if one or more Proxy/Servers fail. Using VRRP, the Client is unaware which of the (redundant) FHS Proxy/Servers is currently providing service, and any service discontinuity will be limited to the failover time supported by VRRP. Widely deployed public domain implementations of VRRP are available.

Proxy/Servers SHOULD use high availability clustering services so that multiple redundant systems can provide coordinated response to failures. As with VRRP, widely deployed public domain implementations of high availability clustering services are available. Note that special-purpose and expensive dedicated hardware is not necessary, and public domain implementations can be used even between lightweight virtual machines in cloud deployments.

16. Detecting and Responding to Proxy/Server Failures

In environments where fast recovery from Proxy/Server failure is essential, FHS Proxy/Servers SHOULD use proactive Neighbor Unreachability Detection (NUD) in a manner that parallels Bidirectional Forwarding Detection (BFD) [RFC5880] to track MAP Proxy/Server reachability. FHS Proxy/Servers can then quickly detect and react to failures so that cached information is re-established through alternate paths. Proactive NUD control messaging is carried only over well-connected ground domain networks (i.e., and not low-end links such as aeronautical radios) and can therefore be tuned for rapid response.

FHS Proxy/Servers perform proactive NUD for MAP Proxy/Servers for which there are currently active Clients. If a MAP Proxy/Server fails, the FHS Proxy/Server can quickly inform Clients of the outage by sending multicast RA messages. The FHS Proxy/Server sends RA messages to Clients with Source Address set to the GUA of the MAP, with Destination Address set to link-scoped All-Nodes multicast (ff02::1) [RFC4291] and with Router Lifetime set to 0.

The FHS Proxy/Server SHOULD send MAX_FINAL_RTR_ADVERTISEMENTS RA messages separated by small delays [RFC4861]. Any Clients that have been using the (now defunct) MAP Proxy/Server will receive the RA messages.

17. Transition Considerations

When a Client connects to a *NET link for the first time, it sends an RS message with an OMNI option. If the first hop router recognizes the option, it responds according to the appropriate FHS/MAP Proxy/Server role resulting in an RA message with an OMNI option returned to the Client. The Client then engages this FHS Proxy/Server according to the OMNI link model specified above. If the first hop router is a legacy IPv6 router, however, it instead returns an RA message with no OMNI option and with an ordinary unicast source LLA as specified in [RFC4861]. In that case, the Client engages the *NET according to the legacy IPv6 link model and without the OMNI extensions specified in this document.

If the *NET link model is multiple access, there must be assurance that address duplication cannot corrupt the neighbor caches of other nodes on the link. When the Client sends an RS message on a multiple access *NET link with an OMNI option, first hop routers that recognize the option ensure that the Client is authorized to use the address and return an RA with a non-zero Router Lifetime only if the Client is authorized. First hop routers that do not recognize the OMNI option instead return an RA that makes no statement about the Client's authorization to use the Source Address. In that case, the Client should perform Duplicate Address Detection to ensure that it does not interfere with other nodes on the link.

An alternative approach for multiple access *NET links to ensure isolation for Client-Proxy/Server communications is through link layer address mappings as discussed in Appendix D. This arrangement imparts a (virtual) point-to-point link model over the (physical) multiple access link.

18. OMNI Interfaces on Open Internetworks

Client OMNI interfaces configured over IPv6-enabled underlay interfaces on an open Internetwork without an OMNI-aware first-hop router receive IPv6 RA messages with no OMNI options, while OMNI interfaces configured over IPv4-only underlay interfaces receive no IPv6 RA messages at all (but may receive IPv4 RA messages per [RFC1256]). Client OMNI interfaces that receive RA messages with OMNI options configure addresses, on-link prefixes, etc. on the underlay interface that received the RA according to standard IPv6 ND and address resolution conventions [RFC4861] [RFC4862]. Client OMNI interfaces configured over IPv4-only underlay interfaces configure IPv4 address information on the underlay interfaces using mechanisms such as DHCPv4 [RFC2131].

Client OMNI interfaces configured over underlay interfaces connected to open Internetworks can apply lower layer security services such as VPNs (e.g., IPsec tunnels) to connect to a Proxy/Server, or can establish a secured direct point-to-point link to the Proxy/Server through some other means (see Section 4). In environments where lower layer security may be impractical or undesirable, Client OMNI interfaces can instead send IPv6 ND messages with OMNI sub-options that include authentication parameters.

OMNI interfaces use UDP/IP as L2 encapsulation headers for transmission over open Internetworks with UDP service port number 8060 for both IPv4 and IPv6 underlay interfaces. The OMNI interface submits original IP packets for OAL encapsulation, then encapsulates the resulting OAL fragments in UDP/IP L2 headers to form carrier packets. (The first 4 bits following the UDP header determine

whether the OAL headers are uncompressed/compressed as discussed in Section 6.5.) The OMNI interface sets the UDP length to the encapsulated OAL fragment length and sets the IP length to an appropriate value at least as large as the UDP datagram.

When necessary, sources include an OMNI option with an RSA Signature or HMAC sub-option in IPv6 ND messages. Procedures for including OMNI authentication sub-options are discussed in Section 10.

After establishing a secured underlay link or preparing for UDP/IP encapsulation, OMNI interfaces send RS/RA messages for Client-Proxy/Server coordination (see: Section 13) and peer-to-peer IPv6 ND solicitation and response messages for multilink forwarding, route optimization, and mobility management (see: [I-D.templin-6man-aero3]). These control plane messages must be authenticated while other control and data plane messages are delivered the same as for ordinary best effort traffic with Source Address and/or Identification window-based data origin verification. Transport and higher layer protocol sessions over OMNI interfaces that connect over open Internetworks without an explicit underlay link security services should therefore employ security at their layers to ensure authentication, integrity and/or confidentiality.

Clients should avoid using INET Proxy/Servers as general-purpose routers for steady streams of carrier packets that do not require authentication. Clients should therefore perform route optimization to coordinate with other INET nodes that can provide forwarding services (or preferably coordinate with peer Clients directly) instead of burdening the Proxy/Server. Procedures for coordinating with peer Clients and discovering INET nodes that can provide better forwarding services are discussed in [I-D.templin-6man-aero3].

Clients that attempt to contact peers over INET underlay interfaces often encounter NATs in the path. OMNI interfaces accommodate NAT traversal using UDP/IP encapsulation and the mechanisms discussed in [I-D.templin-6man-aero3]. FHS Proxy/Servers include L2ADDR information in an Interface Attributes sub-option in RA messages to allow Clients to detect the presence of NATs.

Note: Following the initial IPv6 ND message exchange, OMNI interfaces configured over INET underlay interfaces maintain neighbor relationships by transmitting periodic IPv6 ND messages with OMNI options that include authentication signatures.

Note: OMNI interfaces configured over INET underlay interfaces should employ the Identification window synchronization mechanisms specified in Section 6.7 in order to exclude spurious carrier packets that might otherwise clutter the reassembly cache. This is especially important in environments where carrier packet spoofing and/or corruption is a threat.

Note: NATs may be present on the path from a Client to its FHS Proxy/Server, but never on the path from the FHS Proxy/Server to the MAP where only INET and/or spanning tree hops occur. Therefore, the FHS Proxy/Server does not communicate Client origin information to the MAP where it would serve no purpose.

19. Time-Varying MNPs

In some use cases, it is desirable, beneficial and efficient for the Client to receive a constant MNP that travels with the Client wherever it moves. For example, this would allow air traffic controllers to easily track aircraft, etc. In other cases, however (e.g., intelligent transportation systems), the Client may be willing to sacrifice a modicum of efficiency in order to have time-varying MNPs that can be changed occasionally to defeat adversarial tracking.

The prefix delegation services discussed in Section 13.3 allows Clients that desire time-varying MNPs to obtain short-lived prefixes to send RS messages with an OMNI option with DHCPv6 IA_PD sub-options. The Client would then be obligated to renumber its internal networks whenever its MNPs change. This should not present a challenge for Clients with automated network renumbering services, but may disrupt persistent sessions that would prefer to use a constant address.

20. Error Messages

An OAL destination or intermediate system may need to return ICMPv6-like error messages (e.g., Destination Unreachable, Packet Too Big, Time Exceeded, etc.) [RFC4443] to an OAL source. Since ICMPv6 error messages do not themselves include authentication codes, OAL nodes can instead return error messages as an OMNI ICMPv6 Error sub-option in a secured unsolicited IPv6 ND control message.

21. IANA Considerations

The following IANA actions are requested in accordance with [RFC8126]. Both existing registries and new registries specific to OMNI are affected. Existing registries should be updated according to standard IANA practices. New registries should be created under a new registry group for "Overlay Multilink Network (OMNI) Interface".

21.1. Protocol Numbers

The IANA is instructed to allocate an Internet Protocol number TBD1 from the <https://www.iana.org/assignments/protocol-numbers> registry for the Overlay Multilink Network (OMNI) Interface as a non IPv6 Extension Header protocol. Guidance is found in [RFC5237] (registration procedure is IESG Approval or Standards Action).

21.2. IEEE 802 Numbers

During final publication stages, the IESG will be requested to procure an IEEE EtherType value TBD2 for OMNI according to the statement found at <https://www.ietf.org/about/groups/iesg/statements/ethertypes/>.

Following this procurement, the IANA is instructed to register the value TBD2 in the Ethertypes registry of the <https://www.iana.org/assignments/ieee-802-numbers> registry group for "Overlay Multilink Network (OMNI) Interface encapsulation on Ethernet networks". Guidance is found in [RFC7042] (registration procedure is Expert Review).

21.3. IPv4 Special-Purpose Address

The IANA is instructed to assign TBD3/N as an "OMNI IPv4 anycast" address/prefix in the <https://www.iana.org/assignments/iana-ipv4-special-registry> registry in a similar fashion as for [RFC3068]. The assignment also automatically provides the basis for an OMNI IPv6 subnet router anycast address configured as 2002:TBD3::. The IANA is requested to assist the author's efforts to obtain a TBD3/N public IPv4 prefix, whether through an RIR allocation, a delegation from the "Current Recovered IPv4 Pool" of the <https://www.iana.org/assignments/ipv4-recovered-address-space> registry or through an unspecified third party donation.

21.4. IANA OUI Ethernet Numbers

The IANA is instructed to allocate one Ethernet unicast address TBD4 (suggested value '00-52-14') in the "IANA Unicast 48-bit MAC Addresses" registry in the <https://www.iana.org/assignments/ethernet-numbers> registry group (registration procedure is Expert Review). The registration should appear as follows:

Addresses	Usage	Reference
-----	-----	-----
00-52-14	Overlay Multilink Network (OMNI) Interface	[RFCXXXX]

Figure 39: IANA Unicast 48-bit MAC Addresses

21.5. Overlay Multilink Network (OMNI) Interface Registry Group

The IANA is instructed to create a new 'omni-interface' registry group for "Overlay Multilink Network (OMNI) Interface". The registry group must include the following new registries:

21.5.1. OMNI Option Sub-Types (New Registry)

The OMNI option defines a 5-bit Sub-Type field, for which IANA is instructed to create and maintain a new registry entitled "OMNI Option Sub-Type Values". Initial values are given below (registration procedure is RFC required):

Value	Sub-Type name	Reference
-----	-----	-----
0	Pad1	[RFCXXXX]
1	PadN	[RFCXXXX]
2	Node Identification	[RFCXXXX]
3	CGA	[RFCXXXX]
4	RSA Signature	[RFCXXXX]
5	Timestamp	[RFCXXXX]
6	Nonce	[RFCXXXX]
7	Trust Anchor	[RFCXXXX]
8	Certificate	[RFCXXXX]
9	HMAC	[RFCXXXX]
10	Neighbor Synchronization	[RFCXXXX]
11	Interface Attributes	[RFCXXXX]
12	Traffic Selector	[RFCXXXX]
13	Geo Coordinates	[RFCXXXX]
14	DHCPv6 Message	[RFCXXXX]
15	PIM-SM Message	[RFCXXXX]
16	Fragmentation Report	[RFCXXXX]
17	ICMPv6 Error	[RFCXXXX]
18	Proxy/Server Departure	[RFCXXXX]
19-252	Unassigned	
253-254	Reserved for Experimentation	[RFCXXXX]
255	Reserved by IANA	[RFCXXXX]

Figure 40: OMNI Option Sub-Type Values

21.5.2. OMNI Node Identification ID-Types (New Registry)

The OMNI Node Identification sub-option (see: Section 10.2.3) contains an 8-bit ID-Type field, for which IANA is instructed to create and maintain a new registry entitled "OMNI Node Identification ID-Type Values". Initial values are given below (registration procedure is RFC required):

Value	Sub-Type name	Reference
-----	-----	-----
0	MLA	[RFCXXXX]
1	UUID	[RFCXXXX]
2	Network Access Identifier	[RFCXXXX]
3	FQDN	[RFCXXXX]
4	IPv4 Address	[RFCXXXX]
5	Unassigned	[RFCXXXX]
6	IPv6 Address	[RFCXXXX]
7-65532	Unassigned	[RFCXXXX]
65533	Reserved for Experimental Use	[RFCXXXX]
65534	Reserved for Experimental Use	[RFCXXXX]
65535	Reserved by IANA	[RFCXXXX]

Figure 41: OMNI Node Identification ID-Type Values

21.5.3. OMNI Geo Coordinates Types (New Registry)

The OMNI Geo Coordinates sub-option (see: Section 10.2.14) contains an 8-bit Type field, for which IANA is instructed to create and maintain a new registry entitled "OMNI Geo Coordinates Type Values". Initial values are given below (registration procedure is RFC required):

Value	Sub-Type name	Reference
-----	-----	-----
0	NULL	[RFCXXXX]
1-252	Unassigned	[RFCXXXX]
253-254	Reserved for Experimental Use	[RFCXXXX]
255	Reserved by IANA	[RFCXXXX]

Figure 42: OMNI Geo Coordinates Type

21.6. Additional Considerations

IANA has assigned UDP port number "8060" for an earlier experimental version of AERO [RFC6706]. This document reclaims UDP port number "8060" for 'aero' as the service port for OMNI interface UDP/IP encapsulation. (Note that, although [RFC6706] is not widely implemented or deployed, any messages coded to that specification can be easily distinguished and ignored since they include an invalid ICMPv6 message type number '0'.) IANA is therefore instructed to update the reference for UDP port number "8060" from "RFC6706" to "RFCXXXX" (i.e., this document) while retaining the existing name 'aero'.

IANA has assigned a 4-octet Private Enterprise Number (PEN) code "45282" in the <https://www.iana.org/assignments/enterprise-numbers> registry. This document is the normative reference for using code "45282" in DHCP Unique IDentifiers based on Enterprise Numbers ("DUID-EN for OMNI Interfaces") (see: Section 9). IANA is therefore instructed to change the enterprise designation for PEN code "45282" from "LinkUp Networks" to "Overlay Multilink Network (OMNI) Interface".

IANA has assigned the ifType code "301 - omni - Overlay Multilink Network (OMNI) Interface" in accordance with Section 6 of [RFC8892]. The registration appears under the IANA <https://www.iana.org/assignments/smi-numbers> registry group Interface Types (ifType)" registry, and the IANA is instructed to change the reference to [RFCXXXX] (i.e., this document).

No further IANA actions are required.

22. Security Considerations

Security considerations for IPv4 [RFC0791], IPv6 [RFC8200] and IPv6 Neighbor Discovery [RFC4861] apply. For end-to-end peer exchanges not fully protected by security associations, OMNI interfaces SHOULD use SECure Neighbor Discovery (SEND) per [RFC3971] or a Hashed Message Authentication Code (HMAC) per [RFC8754] as an adaptation layer service to ensure authentic exchanges. These services provide authentication for unsecured FHS and LHS *NETs, while intermediate hops are protected by the secured spanning tree (see below).

OMNI interfaces configured over secured ANET/ENET interfaces inherit the physical and/or link layer security properties (i.e., "protected spectrum") of the connected networks. OMNI interfaces configured over open *NET interfaces can use symmetric securing services such as IPsec tunnels [RFC4301] or can by some other means establish a direct point-to-point link secured at lower layers.

OMNI link mobility services MUST support strong authentication for control plane messages and forwarding path integrity for data plane messages. In particular, the AERO service [I-D.templin-6man-aero3] constructs a secured spanning tree with Proxy/Servers as leaf nodes and secures the spanning tree links with network layer security services based on IPsec [RFC4301] with IKEv2 [RFC7296]. (Note that direct point-to-point links secured at lower layers can also be used instead of or in addition to network layer security.) Together, these services provide connectionless integrity and data origin authentication with optional protection against replays.

Control plane messages that affect the routing system or neighbor state either include authentication signatures or are constrained to travel only over secured spanning tree paths; in both cases, the messages are protected by network (and/or lower-layer) security. Other control and data plane messages can travel over unsecured route optimized paths that do not strictly follow the spanning tree, therefore end-to-end sessions should employ transport or higher layer security services (e.g., TLS/SSL [RFC8446], DTLS [RFC6347], etc.). Additionally, the OAL Identification value can provide a first level of data origin authentication to mitigate off-path spoofing.

Identity-based key verification infrastructure services such as iPSK may be necessary for verifying the identities claimed by Clients. This requirement should be harmonized with the manner in which identifiers such as MLAs are certified in a given operational environment.

Security considerations for specific access network interface types are covered under the corresponding IP-over-(foo) specification (e.g., [RFC2464], [RFC2492], etc.).

Security considerations for IPv6 fragmentation and reassembly are discussed in Section 6.13. In environments where spoofing is considered a threat, OMNI nodes SHOULD employ Identification window synchronization and OAL destinations SHOULD configure an (end-system-based) firewall.

23. Implementation Status

AERO/OMNI Release-3.2 was tagged on March 30, 2021, and was subject to internal testing. The implementation is not planned for public release.

A write-from-scratch reference implementation is under active internal development, with release version v0.1 tagged on December 9, 2024 and version v0.2 tagged on January 22, 2025. Future versions will be made available for public release.

24. Document Updates

This document suggests that the following could be updated through future IETF initiatives:

- * [RFC1191]
- * [RFC4443]
- * [RFC8200]

* [RFC8201]

Updates can be through, e.g., standards action, the errata process, etc. as appropriate.

25. Acknowledgements

The first version of this document was prepared per the consensus decision at the 7th Conference of the International Civil Aviation Organization (ICAO) Working Group-I Mobility Subgroup on March 22, 2019. Consensus to take the document forward to the IETF was reached at the 9th Conference of the Mobility Subgroup on November 22, 2019. Attendees and contributors included: Guray Acar, Danny Bharj, Francois D'Humieres, Pavel Drasil, Nikos Fistas, Giovanni Garofolo, Bernhard Haindl, Vaughn Maiolla, Tom McParland, Victor Moreno, Madhu Niraula, Brent Phillips, Liviu Popescu, Jacky Pouzet, Alope Roy, Greg Saccone, Robert Segers, Michal Skorepa, Michel Solery, Stephane Tamalet, Fred Templin, Jean-Marc Vacher, Bela Varkonyi, Tony Whyman, Fryderyk Wrobel and Dongsong Zeng.

The following individuals are acknowledged for their useful comments: Felipe Magno de Almeida, Amanda Baber, Scott Burleigh, Stuart Card, Donald Eastlake, Adrian Farrel, Michael Matyas, Robert Moskowitz, Madhu Niraula, Greg Saccone, Stephane Tamalet, Eliot Lear, Eduard Vasilenko, Eric Vyncke. Pavel Drasil, Zdenek Jaron and Michal Skorepa are especially recognized for their many helpful ideas and suggestions. Akash Agarwal, Madhuri Madhava Badgandi, Sean Dickson, Don Dillenburg, Joe Dudkowski, Vijayasathy Rajagopalan, Ron Sackman, Bhargava Raman Sai Prakash and Katherine Tran are acknowledged for their hard work on the implementation and technical insights that led to improvements for the spec.

Discussions on the IETF 6man and atn mailing lists during the fall of 2020 suggested additional points to consider. The authors gratefully acknowledge the list members who contributed valuable insights through those discussions. Eric Vyncke and Erik Kline were the intarea ADs, while Bob Hinden and Ole Troan were the 6man WG chairs at the time the document was developed; they are all gratefully acknowledged for their many helpful insights. Many of the ideas in this document have further built on IETF experiences beginning in the 1990s, with insights from colleagues including Ron Bonica, Brian Carpenter, Ralph Droms, Tom Herbert, Bob Hinden, Christian Huitema, Thomas Narten, Dave Thaler, Joe Touch, Pascal Thubert, and many others who deserve recognition.

Early observations on IP fragmentation performance implications were noted in the 1986 Digital Equipment Corporation (DEC) "qe reset" investigation, where fragment bursts from NFS UDP traffic triggered

hardware resets resulting in communication failures. Jeff Chase, Fred Glover and Chet Juzsaszak of the Ultrix Engineering Group led the investigation, and determined that setting a smaller NFS mount block size reduced the amount of fragmentation and suppressed the resets. Early observations on L2 media MTU issues were noted in the 1988 DEC FDDI investigation, where Raj Jain, KK Ramakrishnan and Kathy Wilde represented architectural considerations for FDDI networking in general including FDDI/Ethernet bridging. Jeff Mogul (who led the IETF Path MTU Discovery working group) and other DEC colleagues who supported these early investigations are also acknowledged.

Throughout the 1990's and into the 2000's, many colleagues supported and encouraged continuation of the work. Beginning with the DEC Project Sequoia effort at the University of California, Berkeley, then moving to the DEC research lab offices in Palo Alto CA, then to Sterling Software at the NASA Ames Research Center, then to SRI in Menlo Park, CA, then to Nokia in Mountain View, CA and finally to the Boeing Company in 2005 the work saw continuous advancement through the encouragement of many. Those who offered their support and encouragement are gratefully acknowledged.

This work is aligned with the NASA Safe Autonomous Systems Operation (SASO) program under NASA contract number NNA16BD84C.

This work is aligned with the FAA as per the SE2025 contract number DTFAWA-15-D-00030.

This work is aligned with the Boeing Information Technology (BIT) Mobility Vision Lab (MVL) program.

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

Honoring life, liberty and the pursuit of happiness.

26. References

26.1. Normative References

[I-D.ietf-dhc-rfc8415bis]

Mrugalski, T., Volz, B., Richardson, M., Jiang, S., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", Work in Progress, Internet-Draft, draft-ietf-dhc-rfc8415bis-12, 4 June 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-dhc-rfc8415bis-12>>.

- [I-D.templin-6man-ipid-ext2]
Templin, F. and T. Herbert, "IPv6 Extended Fragment Header (EFH)", Work in Progress, Internet-Draft, draft-templin-6man-ipid-ext2-17, 3 June 2025,
<<https://datatracker.ietf.org/doc/html/draft-templin-6man-ipid-ext2-17>>.
- [RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980,
<<https://www.rfc-editor.org/info/rfc768>>.
- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981,
<<https://www.rfc-editor.org/info/rfc791>>.
- [RFC2003] Perkins, C., "IP Encapsulation within IP", RFC 2003, DOI 10.17487/RFC2003, October 1996,
<<https://www.rfc-editor.org/info/rfc2003>>.
- [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>>.
- [RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, DOI 10.17487/RFC2473, December 1998, <<https://www.rfc-editor.org/info/rfc2473>>.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005,
<<https://www.rfc-editor.org/info/rfc3971>>.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, DOI 10.17487/RFC3972, March 2005,
<<https://www.rfc-editor.org/info/rfc3972>>.
- [RFC4007] Deering, S., Haberman, B., Jinmei, T., Nordmark, E., and B. Zill, "IPv6 Scoped Address Architecture", RFC 4007, DOI 10.17487/RFC4007, March 2005,
<<https://www.rfc-editor.org/info/rfc4007>>.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005,
<<https://www.rfc-editor.org/info/rfc4193>>.

- [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>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [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>>.
- [RFC6088] Tsirtsis, G., Giarreta, G., Soliman, H., and N. Montavont, "Traffic Selectors for Flow Bindings", RFC 6088, DOI 10.17487/RFC6088, January 2011, <<https://www.rfc-editor.org/info/rfc6088>>.
- [RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, DOI 10.17487/RFC6437, November 2011, <<https://www.rfc-editor.org/info/rfc6437>>.
- [RFC6438] Carpenter, B. and S. Amante, "Using the IPv6 Flow Label for Equal Cost Multipath Routing and Link Aggregation in Tunnels", RFC 6438, DOI 10.17487/RFC6438, November 2011, <<https://www.rfc-editor.org/info/rfc6438>>.
- [RFC8028] Baker, F. and B. Carpenter, "First-Hop Router Selection by Hosts in a Multi-Prefix Network", RFC 8028, DOI 10.17487/RFC8028, November 2016, <<https://www.rfc-editor.org/info/rfc8028>>.
- [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>>.
- [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>>.

- [RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed., "Path MTU Discovery for IP version 6", STD 87, RFC 8201, DOI 10.17487/RFC8201, July 2017, <<https://www.rfc-editor.org/info/rfc8201>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC9293] Eddy, W., Ed., "Transmission Control Protocol (TCP)", STD 7, RFC 9293, DOI 10.17487/RFC9293, August 2022, <<https://www.rfc-editor.org/info/rfc9293>>.

26.2. Informative References

- [ATN] Maiolla, V., "The OMNI Interface - An IPv6 Air/Ground Interface for Civil Aviation, IETF Liaison Statement #1676, <https://datatracker.ietf.org/liaison/1676/>", 3 March 2020.
- [ATN-IPS] "ICAO Document 9896 (Manual on the Aeronautical Telecommunication Network (ATN) using Internet Protocol Suite (IPS) Standards and Protocol), Draft Edition 3 (work-in-progress)", 10 December 2020.
- [CRC] Jain, R., "Error Characteristics of Fiber Distributed Data Interface (FDDI), IEEE Transactions on Communications", August 1990.
- [EUI] "IEEE Guidelines for Use of Extended Unique Identifier (EUI), Organizationally Unique Identifier (OUI), and Company ID, <https://standards.ieee.org/wp-content/uploads/import/documents/tutorials/eui.pdf>", 3 August 2017.
- [I-D.gont-dhcbwg-dhcpv6-iids]
Gont, F., "A Method for Generating Semantically Opaque IPv6 Interface Identifiers (IIDs) with Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", Work in Progress, Internet-Draft, draft-gont-dhcbwg-dhcpv6-iids-00, 10 February 2025, <<https://datatracker.ietf.org/doc/html/draft-gont-dhcbwg-dhcpv6-iids-00>>.

[I-D.herbert-ipv4-eh]

Herbert, T., "IPv4 Extension Headers and Flow Label", Work in Progress, Internet-Draft, draft-herbert-ipv4-eh-03, 22 February 2024, <<https://datatracker.ietf.org/doc/html/draft-herbert-ipv4-eh-03>>.

[I-D.ietf-6man-eh-limits]

Herbert, T., "Limits on Sending and Processing IPv6 Extension Headers", Work in Progress, Internet-Draft, draft-ietf-6man-eh-limits-19, 27 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-6man-eh-limits-19>>.

[I-D.ietf-6man-rfc6724-update]

Buraglio, N., Chown, T., and J. Duncan, "Prioritizing known-local IPv6 ULAs through address selection policy", Work in Progress, Internet-Draft, draft-ietf-6man-rfc6724-update-20, 13 May 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-6man-rfc6724-update-20>>.

[I-D.ietf-intarea-tunnels]

Touch, J. D. and M. Townsley, "IP Tunnels in the Internet Architecture", Work in Progress, Internet-Draft, draft-ietf-intarea-tunnels-15, 9 May 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-intarea-tunnels-15>>.

[I-D.ietf-tsvwg-udp-options]

Touch, J. D. and C. M. Heard, "Transport Options for UDP", Work in Progress, Internet-Draft, draft-ietf-tsvwg-udp-options-45, 16 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-tsvwg-udp-options-45>>.

[I-D.ietf-v6ops-ula-usage-considerations]

Jiang, S., Liu, B., and N. Buraglio, "Considerations For Using Unique Local Addresses", Work in Progress, Internet-Draft, draft-ietf-v6ops-ula-usage-considerations-05, 11 December 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-v6ops-ula-usage-considerations-05>>.

[I-D.link-6man-gulla]

Linkova, J., "Using Prefix-Specific Link-Local Addresses to Improve SLAAC Robustness", Work in Progress, Internet-Draft, draft-link-6man-gulla-01, 3 March 2025, <<https://datatracker.ietf.org/doc/html/draft-link-6man-gulla-01>>.

[I-D.perkins-manet-aodvv2]

Perkins, C. E., Dowdell, J., Steenbrink, L., and V. Pritchard, "Ad Hoc On-demand Distance Vector Version 2 (AODVv2) Routing", Work in Progress, Internet-Draft, draft-perkins-manet-aodvv2-05, 22 November 2024, <<https://datatracker.ietf.org/doc/html/draft-perkins-manet-aodvv2-05>>.

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

Templin, F. L., "Automatic Extended Route Optimization (AERO)", Work in Progress, Internet-Draft, draft-templin-6man-aero3-45, 13 June 2025, <<https://datatracker.ietf.org/doc/html/draft-templin-6man-aero3-45>>.

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

Templin, F., "IPv6 Addresses for Ad Hoc Networks", Work in Progress, Internet-Draft, draft-templin-6man-mla-27, 3 April 2025, <<https://datatracker.ietf.org/doc/html/draft-templin-6man-mla-27>>.

[I-D.templin-6man-parcels2]

Templin, F., "IPv6 Parcels and Advanced Jumbos (AJs)", Work in Progress, Internet-Draft, draft-templin-6man-parcels2-27, 21 May 2025, <<https://datatracker.ietf.org/doc/html/draft-templin-6man-parcels2-27>>.

[I-D.templin-intarea-parcels2]

Templin, F., "IPv4 Parcels and Advanced Jumbos (AJs)", Work in Progress, Internet-Draft, draft-templin-intarea-parcels2-18, 21 May 2025, <<https://datatracker.ietf.org/doc/html/draft-templin-intarea-parcels2-18>>.

[IANA-CGA] Postel, J., "Cryptographically Generated Addresses (CGA) Message Type Name Space, <https://www.iana.org/assignments/cga-message-types/cga-message-types.xhtml>", 15 March 2023.

[IEEE802.1AX]

"Institute of Electrical and Electronics Engineers, Link Aggregation, IEEE Standard 802.1AX-2008, <https://standards.ieee.org/ieee/802.1AX/6768/>", 29 May 2020.

[IPV4]

Postel, J., "IPv4 Address Space Registry, <https://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml>", 14 December 2020.

- [IPV6] Postel, J., "IPv6 Global Unicast Address Assignments, <https://www.iana.org/assignments/ipv6-unicast-address-assignments/ipv6-unicast-address-assignments.xhtml>", 14 December 2020.
- [RFC0863] Postel, J., "Discard Protocol", STD 21, RFC 863, DOI 10.17487/RFC0863, May 1983, <<https://www.rfc-editor.org/info/rfc863>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [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>>.
- [RFC1149] Waitzman, D., "Standard for the transmission of IP datagrams on avian carriers", RFC 1149, DOI 10.17487/RFC1149, April 1990, <<https://www.rfc-editor.org/info/rfc1149>>.
- [RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, DOI 10.17487/RFC1191, November 1990, <<https://www.rfc-editor.org/info/rfc1191>>.
- [RFC1256] Deering, S., Ed., "ICMP Router Discovery Messages", RFC 1256, DOI 10.17487/RFC1256, September 1991, <<https://www.rfc-editor.org/info/rfc1256>>.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, DOI 10.17487/RFC2104, February 1997, <<https://www.rfc-editor.org/info/rfc2104>>.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <<https://www.rfc-editor.org/info/rfc2131>>.
- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, DOI 10.17487/RFC2464, December 1998, <<https://www.rfc-editor.org/info/rfc2464>>.

- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <<https://www.rfc-editor.org/info/rfc2474>>.
- [RFC2492] Armitage, G., Schulter, P., and M. Jork, "IPv6 over ATM Networks", RFC 2492, DOI 10.17487/RFC2492, January 1999, <<https://www.rfc-editor.org/info/rfc2492>>.
- [RFC2675] Borman, D., Deering, S., and R. Hinden, "IPv6 Jumbograms", RFC 2675, DOI 10.17487/RFC2675, August 1999, <<https://www.rfc-editor.org/info/rfc2675>>.
- [RFC2863] McCloghrie, K. and F. Kastenholtz, "The Interfaces Group MIB", RFC 2863, DOI 10.17487/RFC2863, June 2000, <<https://www.rfc-editor.org/info/rfc2863>>.
- [RFC2923] Lahey, K., "TCP Problems with Path MTU Discovery", RFC 2923, DOI 10.17487/RFC2923, September 2000, <<https://www.rfc-editor.org/info/rfc2923>>.
- [RFC2983] Black, D., "Differentiated Services and Tunnels", RFC 2983, DOI 10.17487/RFC2983, October 2000, <<https://www.rfc-editor.org/info/rfc2983>>.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, DOI 10.17487/RFC3056, February 2001, <<https://www.rfc-editor.org/info/rfc3056>>.
- [RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, DOI 10.17487/RFC3068, June 2001, <<https://www.rfc-editor.org/info/rfc3068>>.
- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <<https://www.rfc-editor.org/info/rfc3168>>.
- [RFC3366] Fairhurst, G. and L. Wood, "Advice to link designers on link Automatic Repeat reQuest (ARQ)", BCP 62, RFC 3366, DOI 10.17487/RFC3366, August 2002, <<https://www.rfc-editor.org/info/rfc3366>>.
- [RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", BCP 82, RFC 3692, DOI 10.17487/RFC3692, January 2004, <<https://www.rfc-editor.org/info/rfc3692>>.

- [RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<https://www.rfc-editor.org/info/rfc3810>>.
- [RFC3819] Karn, P., Ed., Bormann, C., Fairhurst, G., Grossman, D., Ludwig, R., Mahdavi, J., Montenegro, G., Touch, J., and L. Wood, "Advice for Internet Subnetwork Designers", BCP 89, RFC 3819, DOI 10.17487/RFC3819, July 2004, <<https://www.rfc-editor.org/info/rfc3819>>.
- [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, DOI 10.17487/RFC4213, October 2005, <<https://www.rfc-editor.org/info/rfc4213>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <<https://www.rfc-editor.org/info/rfc4302>>.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <<https://www.rfc-editor.org/info/rfc4303>>.
- [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, DOI 10.17487/RFC4380, February 2006, <<https://www.rfc-editor.org/info/rfc4380>>.
- [RFC4389] Thaler, D., Talwar, M., and C. Patel, "Neighbor Discovery Proxies (ND Proxy)", RFC 4389, DOI 10.17487/RFC4389, April 2006, <<https://www.rfc-editor.org/info/rfc4389>>.
- [RFC4541] Christensen, M., Kimball, K., and F. Solensky, "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", RFC 4541, DOI 10.17487/RFC4541, May 2006, <<https://www.rfc-editor.org/info/rfc4541>>.
- [RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, DOI 10.17487/RFC4605, August 2006, <<https://www.rfc-editor.org/info/rfc4605>>.

- [RFC4963] Heffner, J., Mathis, M., and B. Chandler, "IPv4 Reassembly Errors at High Data Rates", RFC 4963, DOI 10.17487/RFC4963, July 2007, <<https://www.rfc-editor.org/info/rfc4963>>.
- [RFC5213] Gundavelli, S., Ed., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, DOI 10.17487/RFC5213, August 2008, <<https://www.rfc-editor.org/info/rfc5213>>.
- [RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, DOI 10.17487/RFC5214, March 2008, <<https://www.rfc-editor.org/info/rfc5214>>.
- [RFC5237] Arkko, J. and S. Bradner, "IANA Allocation Guidelines for the Protocol Field", BCP 37, RFC 5237, DOI 10.17487/RFC5237, February 2008, <<https://www.rfc-editor.org/info/rfc5237>>.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <<https://www.rfc-editor.org/info/rfc5340>>.
- [RFC5558] Templin, F., Ed., "Virtual Enterprise Traversal (VET)", RFC 5558, DOI 10.17487/RFC5558, February 2010, <<https://www.rfc-editor.org/info/rfc5558>>.
- [RFC5614] Ogier, R. and P. Spagnolo, "Mobile Ad Hoc Network (MANET) Extension of OSPF Using Connected Dominating Set (CDS) Flooding", RFC 5614, DOI 10.17487/RFC5614, August 2009, <<https://www.rfc-editor.org/info/rfc5614>>.
- [RFC5798] Nadas, S., Ed., "Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6", RFC 5798, DOI 10.17487/RFC5798, March 2010, <<https://www.rfc-editor.org/info/rfc5798>>.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.
- [RFC5889] Baccelli, E., Ed. and M. Townsley, Ed., "IP Addressing Model in Ad Hoc Networks", RFC 5889, DOI 10.17487/RFC5889, September 2010, <<https://www.rfc-editor.org/info/rfc5889>>.

- [RFC5942] Singh, H., Beebee, W., and E. Nordmark, "IPv6 Subnet Model: The Relationship between Links and Subnet Prefixes", RFC 5942, DOI 10.17487/RFC5942, July 2010, <<https://www.rfc-editor.org/info/rfc5942>>.
- [RFC6081] Thaler, D., "Teredo Extensions", RFC 6081, DOI 10.17487/RFC6081, January 2011, <<https://www.rfc-editor.org/info/rfc6081>>.
- [RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, DOI 10.17487/RFC6145, April 2011, <<https://www.rfc-editor.org/info/rfc6145>>.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <<https://www.rfc-editor.org/info/rfc6146>>.
- [RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", RFC 6147, DOI 10.17487/RFC6147, April 2011, <<https://www.rfc-editor.org/info/rfc6147>>.
- [RFC6214] Carpenter, B. and R. Hinden, "Adaptation of RFC 1149 for IPv6", RFC 6214, DOI 10.17487/RFC6214, April 2011, <<https://www.rfc-editor.org/info/rfc6214>>.
- [RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, DOI 10.17487/RFC6296, June 2011, <<https://www.rfc-editor.org/info/rfc6296>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [RFC6543] Gundavelli, S., "Reserved IPv6 Interface Identifier for Proxy Mobile IPv6", RFC 6543, DOI 10.17487/RFC6543, May 2012, <<https://www.rfc-editor.org/info/rfc6543>>.
- [RFC6706] Templin, F., Ed., "Asymmetric Extended Route Optimization (AERO)", RFC 6706, DOI 10.17487/RFC6706, August 2012, <<https://www.rfc-editor.org/info/rfc6706>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <<https://www.rfc-editor.org/info/rfc6724>>.

- [RFC6762] Cheshire, S. and M. Krochmal, "Multicast DNS", RFC 6762, DOI 10.17487/RFC6762, February 2013, <<https://www.rfc-editor.org/info/rfc6762>>.
- [RFC6890] Cotton, M., Vegoda, L., Bonica, R., Ed., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, DOI 10.17487/RFC6890, April 2013, <<https://www.rfc-editor.org/info/rfc6890>>.
- [RFC6935] Eubanks, M., Chimento, P., and M. Westerlund, "IPv6 and UDP Checksums for Tunneled Packets", RFC 6935, DOI 10.17487/RFC6935, April 2013, <<https://www.rfc-editor.org/info/rfc6935>>.
- [RFC6936] Fairhurst, G. and M. Westerlund, "Applicability Statement for the Use of IPv6 UDP Datagrams with Zero Checksums", RFC 6936, DOI 10.17487/RFC6936, April 2013, <<https://www.rfc-editor.org/info/rfc6936>>.
- [RFC6980] Gont, F., "Security Implications of IPv6 Fragmentation with IPv6 Neighbor Discovery", RFC 6980, DOI 10.17487/RFC6980, August 2013, <<https://www.rfc-editor.org/info/rfc6980>>.
- [RFC7042] Eastlake 3rd, D. and J. Abley, "IANA Considerations and IETF Protocol and Documentation Usage for IEEE 802 Parameters", RFC 7042, DOI 10.17487/RFC7042, October 2013, <<https://www.rfc-editor.org/info/rfc7042>>.
- [RFC7094] McPherson, D., Oran, D., Thaler, D., and E. Osterweil, "Architectural Considerations of IP Anycast", RFC 7094, DOI 10.17487/RFC7094, January 2014, <<https://www.rfc-editor.org/info/rfc7094>>.
- [RFC7181] Clausen, T., Dearlove, C., Jacquet, P., and U. Herberg, "The Optimized Link State Routing Protocol Version 2", RFC 7181, DOI 10.17487/RFC7181, April 2014, <<https://www.rfc-editor.org/info/rfc7181>>.
- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <<https://www.rfc-editor.org/info/rfc7296>>.
- [RFC7343] Laganier, J. and F. Dupont, "An IPv6 Prefix for Overlay Routable Cryptographic Hash Identifiers Version 2 (ORCHIDv2)", RFC 7343, DOI 10.17487/RFC7343, September 2014, <<https://www.rfc-editor.org/info/rfc7343>>.

- [RFC7542] DeKok, A., "The Network Access Identifier", RFC 7542, DOI 10.17487/RFC7542, May 2015, <<https://www.rfc-editor.org/info/rfc7542>>.
- [RFC7739] Gont, F., "Security Implications of Predictable Fragment Identification Values", RFC 7739, DOI 10.17487/RFC7739, February 2016, <<https://www.rfc-editor.org/info/rfc7739>>.
- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <<https://www.rfc-editor.org/info/rfc7761>>.
- [RFC7847] Melia, T., Ed. and S. Gundavelli, Ed., "Logical-Interface Support for IP Hosts with Multi-Access Support", RFC 7847, DOI 10.17487/RFC7847, May 2016, <<https://www.rfc-editor.org/info/rfc7847>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8402] Filts, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <<https://www.rfc-editor.org/info/rfc8799>>.
- [RFC8892] Thaler, D. and D. Romascanu, "Guidelines and Registration Procedures for Interface Types and Tunnel Types", RFC 8892, DOI 10.17487/RFC8892, August 2020, <<https://www.rfc-editor.org/info/rfc8892>>.
- [RFC8899] Fairhurst, G., Jones, T., T端 xen, M., R端 ngeler, I., and T. V端 lker, "Packetization Layer Path MTU Discovery for Datagram Transports", RFC 8899, DOI 10.17487/RFC8899, September 2020, <<https://www.rfc-editor.org/info/rfc8899>>.

- [RFC8900] Bonica, R., Baker, F., Huston, G., Hinden, R., Troan, O., and F. Gont, "IP Fragmentation Considered Fragile", BCP 230, RFC 8900, DOI 10.17487/RFC8900, September 2020, <<https://www.rfc-editor.org/info/rfc8900>>.
- [RFC9365] Jeong, J., Ed., "IPv6 Wireless Access in Vehicular Environments (IPWAVE): Problem Statement and Use Cases", RFC 9365, DOI 10.17487/RFC9365, March 2023, <<https://www.rfc-editor.org/info/rfc9365>>.
- [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>>.
- [RFC9562] Davis, K., Peabody, B., and P. Leach, "Universally Unique Identifiers (UUIDs)", RFC 9562, DOI 10.17487/RFC9562, May 2024, <<https://www.rfc-editor.org/info/rfc9562>>.
- [RFC9602] Krishnan, S., "Segment Routing over IPv6 (SRv6) Segment Identifiers in the IPv6 Addressing Architecture", RFC 9602, DOI 10.17487/RFC9602, October 2024, <<https://www.rfc-editor.org/info/rfc9602>>.

Appendix A. IPv6 Compatible Addresses

Section 2.5.5.1 of [RFC4291] defines an "IPv4-Compatible IPv6 Address" with the following structure:




Figure 43: IPv4-Compatible IPv6 Address

Although [RFC4291] deprecates the address format from its former use in IPv6 transition mechanisms, this document defines OMNI-specific uses.

When an IPv4-Compatible IPv6 address appears in a packet sent over an OMNI link, the most significant 96 bits are 0 and the least significant 32 bits include an IPv4 address as shown above.

When the address format is used for temporary local address conversions to IPv6, however, it can also be used to represent EUI-48 and EUI-64 addresses as shown below:




Figure 44: EUI-[48/64] Compatible IPv6 Addresses

The above EUI-48 and EUI-64 compatible IPv6 forms MAY be used for temporary local address conversions, such as when converting EUI addresses to IPv6 to support IPv6 fragmentation/reassembly. The address forms MUST NOT appear in the IPv6 headers of packets sent over the OMNI link, however they MAY appear in the body of a packet if also accompanied by a Type designator.

Appendix B. IPv6 ND Message Authentication and Integrity

OMNI interface IPv6 ND messages are subject to authentication and integrity checks at multiple levels. When an OMNI interface sends an IPv6 ND message over an INET interface, it first includes a standard IPv6 ND message checksum, then optionally includes an RSA Signature or HMAC sub-option with a valid signature if necessary then finally always includes an OMNI option checksum. The OMNI interface that receives the message verifies the OMNI option checksum followed by the authentication signature (if present) to ensure IPv6 ND message integrity and authenticity. (The network layer will then verify the IPv6 ND message checksum following OAL processing.)

When an OMNI interface sends an IPv6 ND message over an underlay interface connected to a secured network, it omits the Authentication (sub-)option but always calculates/includes the IPv6 ND message and OMNI option checksums. When an OMNI interface sends an IPv6 ND message over an underlay interface connected to an unsecured network, it first includes an OMNI RSA Signature or HMAC sub-option and calculates the signature beginning with the IPv6 ND message header checksum field and extending to the end of the entire (composite) packet followed by any OMNI sub-options up to but not including the authentication sub-option itself. The OMNI interface next writes the signature into the appropriate field, then calculates the OMNI option checksum as above.

The OMNI interface that receives the message applies any link layer authentication and integrity checks, then verifies the OMNI option checksum. If the checks are correct, the OMNI interface next

verifies the authentication signature. The OMNI interface then delivers the IPv6 ND message to the network layer only if the OMNI option checksum and authentication signature were correct.

OAL destinations also discard carrier packets with unacceptable Identifications and submit the encapsulated fragments in all others for reassembly. The reassembly algorithm rejects any fragments with unacceptable sizes, offsets, etc. and reassembles all others. During reassembly, the extended Identification value provides an integrity assurance vector that complements any integrity checks already applied by lower layers as well as a first-pass filter for any checks that will be applied later by upper layers.

Appendix C. VDL Mode 2 Considerations

ICAO Doc 9776 is the "Technical Manual for VHF Data Link Mode 2" (VDLM2) that specifies an essential radio frequency data link service for aircraft and ground stations in worldwide civil aviation air traffic management. The VDLM2 link type is "multicast capable" [RFC4861], but with considerable differences from common multicast links such as Ethernet and IEEE 802.11.

First, the VDLM2 link data rate is only 31.5Kbps - multiple orders of magnitude less than most modern wireless networking gear. Second, due to the low available link bandwidth only VDLM2 ground stations (i.e., and not aircraft) are permitted to send broadcasts, and even so only as compact link layer "beacons". Third, aircraft employ the services of ground stations by performing unicast RS/RA exchanges upon receipt of beacons instead of listening for multicast RA messages and/or sending multicast RS messages.

This beacon-oriented unicast RS/RA approach is necessary to conserve the already-scarce available link bandwidth. Moreover, since the numbers of beaconing ground stations operating within a given spatial range must be kept as sparse as possible, it would not be feasible to have different classes of ground stations within the same region observing different protocols. It is therefore highly desirable that all ground stations observe a common language of RS/RA as specified in this document.

Note that links of this nature may benefit from compression techniques that reduce the bandwidth necessary for conveying the same amount of data. The IETF lpwan working group is considering possible alternatives: [<https://datatracker.ietf.org/wg/lpwan/documents>].

Appendix D. Client-Proxy/Server Isolation Through Link-Layer Address Mapping

Per [RFC4861], IPv6 ND messages may be sent to either a multicast or unicast link-scoped IPv6 Destination Address. However, IPv6 ND messaging should be coordinated between the Client and Proxy/Server only without invoking other nodes on the underlay network. This implies that Client-Proxy/Server control messaging should be isolated and not overheard by other nodes on the link.

To support Client-Proxy/Server isolation on some links, Proxy/Servers can maintain an OMNI-specific unicast link layer address ("MSADDR"). For Ethernet-compatible links, this specification reserves one Ethernet unicast address TBD4 (see: IANA Considerations). For non-Ethernet statically-addressed links MSADDR is reserved per the assigned numbers authority for the link layer addressing space. For still other links, MSADDR may be dynamically discovered through other means, e.g., link layer beacons.

Clients map the L3 addresses of all IPv6 ND messages they send (i.e., both multicast and unicast) to MSADDR instead of to an ordinary unicast or multicast link layer address. In this way, all of the Client's IPv6 ND messages will be received by Proxy/Servers that are configured to accept carrier packets destined to MSADDR. Note that multiple Proxy/Servers on the link could be configured to accept carrier packets destined to MSADDR, e.g., as a basis for supporting redundancy.

Therefore, Proxy/Servers must accept and process carrier packets destined to MSADDR, while all other devices must not process carrier packets destined to MSADDR. This model has well-established operational experience in Proxy Mobile IPv6 (PMIP) [RFC5213][RFC6543].

Appendix E. IPv6 ND Virtual Interface Model

The OMNI interface linkage between the network and adaptation layers described in this document is based on a virtual Ethernet interface abstraction in a point-to-multipoint configuration. The abstraction allows the network layer and adaptation layer to exchange packets via a virtual Ethernet as though the network layer represents a singular host on one end of the link communicating with a multitude of host entities at the adaptation layer on the other end. This allows the network layer to manage the OMNI interface according to standard IPv6 ND procedures including address resolution, neighbor unreachability detection, duplicate address detection, router discovery and multicast listener discovery.

In an alternative arrangement, the adaptation layer could also emulate a singular host instead of multiple and the virtual link appears as point-to-point. In this model, the network layer configures a static permanent neighbor cache entry for a fictitious hardware address that represents the adaptation layer side of the virtual link. The network layer then forwards all IP packets to this singular adaptation layer neighbor address, and the OMNI interface internally assumes the role of performing all IPv6 ND coordination with external peers without network layer intervention.

While this document is written from the perspective of the point-to-multipoint model, implementations are free to use the point-to-point model as an alternative. Note that it is not required for all nodes on the OMNI link to engage the same model as long as the external appearance of IPv6 ND messages over interconnecting networks is consistent.

Appendix F. IPv4 as an OAL Encapsulation Service

Throughout the document, IPv6 encapsulation per [RFC2473] is assumed as the OMNI Adaptation Layer (OAL) encapsulation service. At first glance, the minimum 40 octets needed for encapsulation may seem excessive however the full OAL encapsulation headers rarely appear over the wire due to effective header compression.

Still, the question may arise as to whether IPv4 encapsulation per [RFC2003] could be applied instead with OMNI encapsulation Type OMNI-IP4. After all, the IPv4 header is significantly smaller than even the smallest IPv6 header plus extensions. However, IPv4 provides only 32-bit addresses useful for OAL addressing whereas IPv6 provides 128-bit addressing allowing for loosely-managed address assignments based on statistical uniqueness.

IPv4 as an OAL encapsulation service may therefore be suitable for small networks where adaptation layer routers operate based on 32-bit router IDs deployed through well-managed assignments. However, IPv4 does not honor the Flow Label and IPv4 links could configure MTUs as small as 68 octets. An OAL IPv4 header plus extensions would also not be as compressible as for IPv6, therefore resulting in extra cost for carrying uncompressible IPv4 header information in the data plane.

Appendix G. Change Log

<< RFC Editor - remove prior to publication >>

Differences from earlier versions:

Draft -60 to -61

- * OMNI sub-option octet alignment requirements clarified.

Draft -59 to -60

- * Updated Interface Attributes Segment List processing procedures; removed "L3ADDR" as it is now part of the Segment List.

Draft -58 to -59

- * Clarification on multi-path Window Synchronization between neighbors.

Draft -57 to -58

- * Removed Parcel Index coding from OCH header formats.
- * Now using raw IPv6 version to indicate presence of an OAL encapsulation header.
- * Included new appendix on IPv4 as an OAL encapsulation.

Draft -56 to -57

- * Globally replaced "AERO Forwarding" with "AERO Flow".
- * Re-ordered AFVI field in OMNI trailer to simplify implementations.

Draft -55 to -56

- * Removed IANA Considerations for ICMP Parameters and ICMPv6 Parameters. The codes are still needed but are now requested in [I-D.templin-6man-ipid-ext2].

Draft -54 to -55

- * Updated Fragmentation Report format and ICMP Error format.
- * Clarifications on "OFS" and Fragmentation Reports.

Draft -52 to -54

- * Removed normative sections on IP parcels and Advanced Jumbos.
- * Revised compressed header format to uniformly use "Res/Index" octets. IP parcels will tell how Parcel ID is coded in these fields.

Draft -51 to -52

- * Support OFS probing based on live data (and not just discard data).

Draft -50 to -51

- * Removed instances of carrier packet source fragmentation.
- * New ICMPv6 Code for "MTU Probe Reply".

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