Transmission of IPv6 Packets over Overlay Multilink Network (OMNI) Interfaces

(draft-templin-6man-omni-interface)

6man WG Interim Session

March 31, 2020

Fred L. Templin (fred.l.templin@boeing.com)

The Boeing Company

Agenda

- Background and Motivation
- Mobile Node / Mobility Service Architecture
- OMNI Interface Model
- MTU
- Link-local Addresses
- OMNI Option
- Router Discovery and Prefix Registration
- Conceptual Multilink Selection Algorithm
- Status and Next Steps

Background and Motivation

- International Civil Aviation Organization (**ICAO**) is designing a worldwide IPv6-based Aeronautical Telecommunications Network with Internet Protocol Services (**ATN/IPS**)
- Each aircraft Mobile Node (**MN**) assigned unique Mobile Network Prefix (**MNP**) from an aggregated Mobility Service Prefix (**MSP**) by the **MSP authority**, e.g., ICAO. For example, /56 MNP can be assigned from /32 MSP - **aircraft as a mobile IPv6 network**
- Aircraft connect to ATN/IPS via multiple air/ground wireless data links often with diverse properties (e.g., VDLM2, LDACS, SATCOM, AeroMACS, others) – **need an Air-to-Ground (A/G) interface for Multilink Coordination**
- Interface specification work began in ICAO, then published as IETF Internet-Draft: **"Transmission of IPv6 Packets over Overlay Multilink Network (OMNI) Interfaces**"
- Substantial review input received resulting in significant improvements; draft re-issued with alignment to 6MAN WG
- Members of ICAO Communications Panel (CP) Working Group I (WG-I) Mobility Subgroup issued "For Action" Liaison Statement 3/29/2020 requesting standards-track publication

Mobile Node / Mobility Service Architecture

- **Mobile Node (MN)** has multiple independent **Access Network (ANET)** data links → **Multilink**
- Native IPv6 (no encapsulation) over ANETs
- **Access Routers (AR)** serve as Proxy between ANET and global **Internetwork (INET)** (e.g., ATN/IPS)
- AR performs encapsulation on its INET interface on behalf of the MN → **Proxy**
- **Mobility Service (MS)** coordinates all mobility and multilinking decisions in the ground-domain INET
- **Mobility Service Endpoints (MSEs)** support **Distributed Mobility Management (DMM)** – each MSE tracks a subset of all MNs in global INET

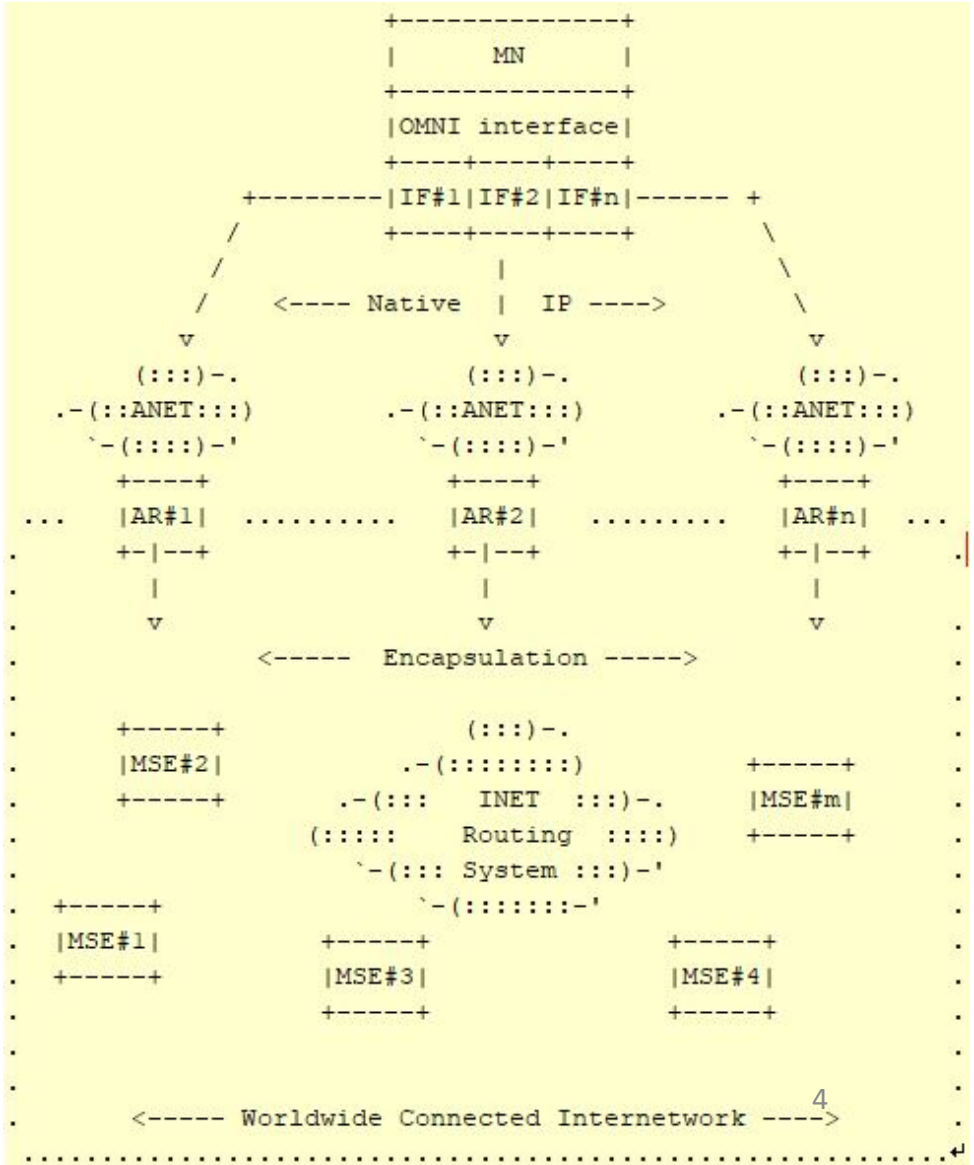

OMNI Interface Model

- **OMNI Interface** is MN's virtual network layer interface configured over multiple diverse underlying ANET interfaces, e.g., VDLM2, SATCOM, LDACS, AeroMACS, etc. **(Multilink)**
- OMNI Interface assigns an IPv6 link-local address known as the **OMNI LLA** - used in **IPv6 Neighbor Discovery (ND)** messaging
- ANET interfaces are **unnumbered** – no IP addresses assigned by the connected ANET
- OMNI Interface is MN's logical attachment to the **OMNI link** - **Non-Broadcast Multiple Access (NBMA)** link model
- OMNI link incorporates the Mobility Service (MS) and spans the entire Internetwork

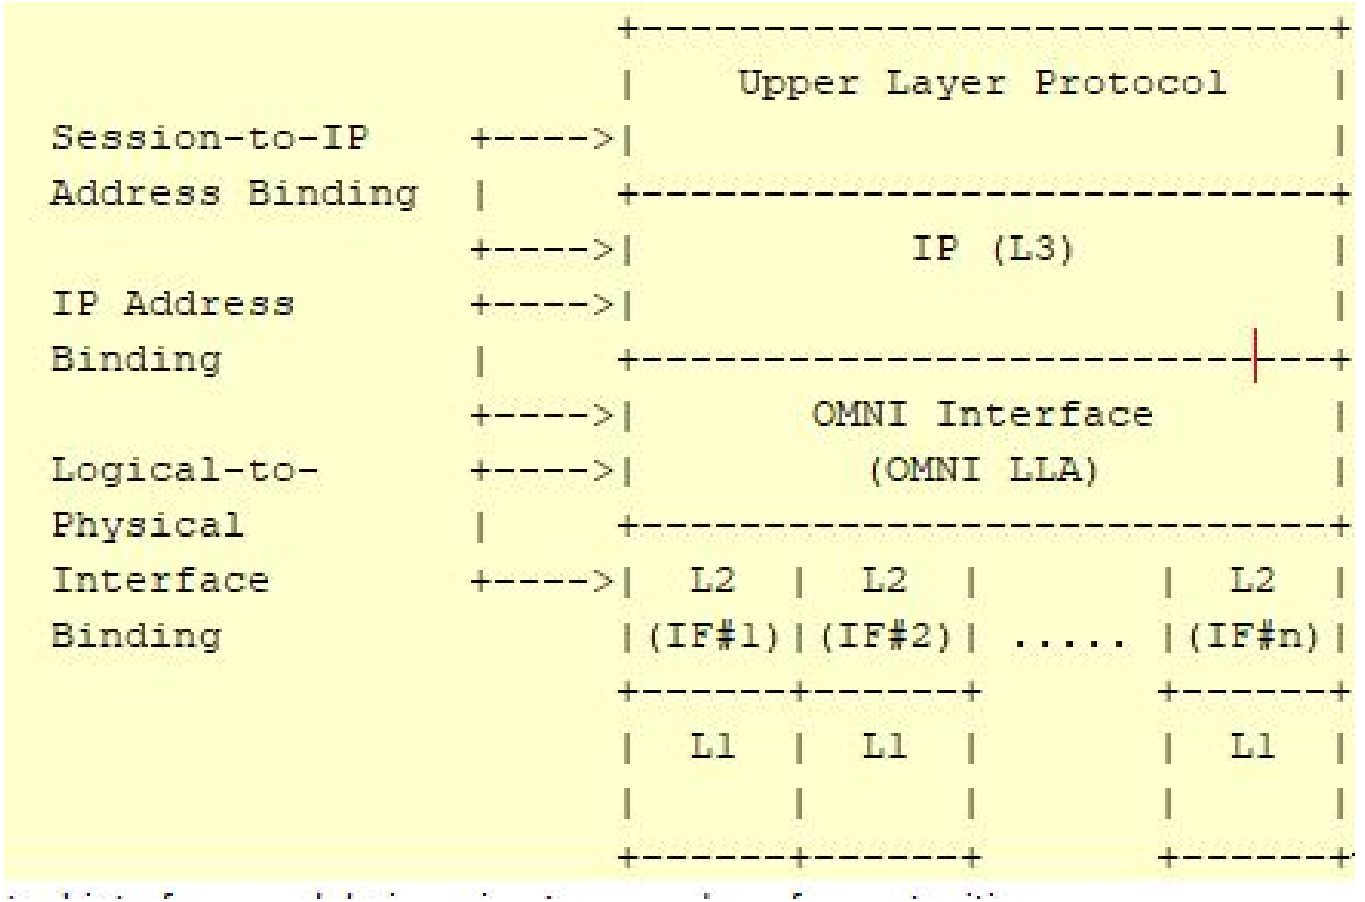

MTU

- **OMNI interface MTU is 9180 [RFC2492]**
- Only the OMNI interface MTU is exposed to the IPv6 layer, while underlying ANET interfaces may have a wide variety of diverse MTUs (e.g., 1280, 1400, 1500, 4K, 9K, etc.)
- Traditional approach:
  - OMNI interface drops packets that are too large for the selected underlying ANET interface and returns internally-generated Packet Too Big (PTB)
- **New approach:**
  - **Within the OMNI interface, perform Link-Local RFC2473 encapsulation**
  - **Apply IPv6 fragmentation to ANET interface MTU size**
  - **Link-local destination reassembles; optionally returns “advisory” PTBs**
  - **Lossless PMTUD – source dynamically adjusts to PMTU with no loss**
  - **Supports sizes up to 9KB GigE Jumbos across any ANET link type**

OMNI Link-Local Addresses (LLAs)

- IPv6 MN OMNI LLAs encode MNP within interface ID of fe80::/64
  - **For the MNP 2001:db8:1000:2000::/56, OMNI LLA is fe80::2001:db8:1000:2000**
- IPv4-Compatible MN OMNI LLAs assigned from fe80::ffff:[v4addr]
  - **For the IPv4 address 192.0.2.1, OMNI LLA is fe80::ffff:192.0.2.1**
- Mobility Service (MS) OMNI LLAs assigned from fe80::/96
  - **fe80::1, fe80::2, fe80::3, … , fe80::feff:ffff**
- IPv6 link-local Subnet-Router anycast address is **fe80::**
- Reserved for future use: **fe80::ff00:0000 – fe80::ffff:ffff**

  - **If IPv6 addressing architecture changes to permit prefixes /65 and longer, OMNI LLA format accommodates up to /118 (future-proof)**

OMNI Option

- Included in OMNI intf. IPv6 ND messages
- New IPv6 ND option - **Type TBD (IANA)**
- **Prefix Length** and **R** used for source address based **Prefix Registration**
  - e.g., for source fe80::2001:db8:1:2, Prefix Length=56, and R=1 means MS registers the MNP 2001:db8:1:2::/56
- A list of sub-options follows with multilink and mobility service parameters

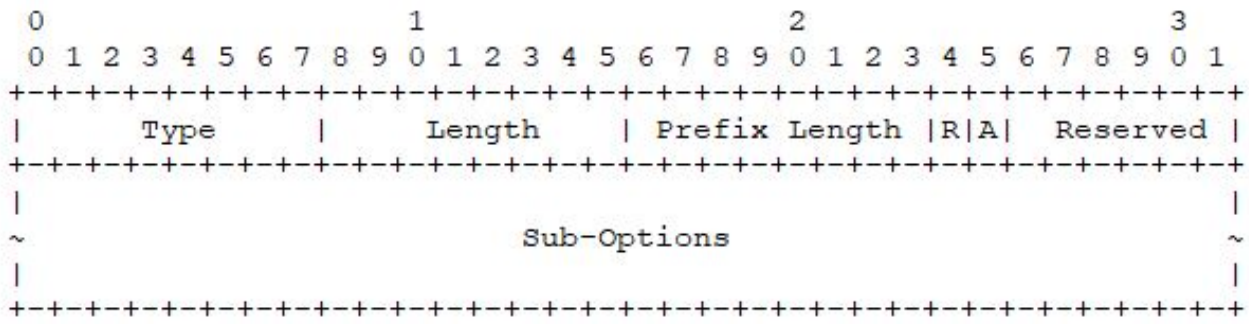

Figure 3: OMNI Option Format

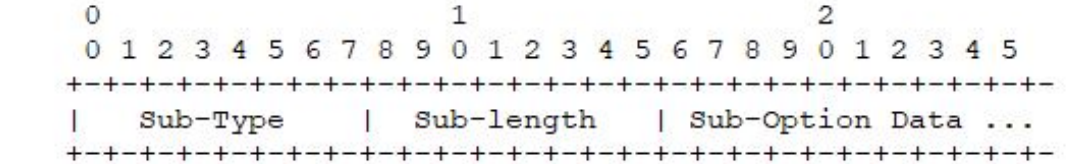

Figure 4: Sub-Option Format

o Sub-Type is a 1-byte field that encodes the Sub-Option type. Sub-Options defined in this document are:

Option Name	Sub-Type
Pad1	0
PadN	1
ifIndex-tuple (Type 1)	2
ifIndex-tuple (Type 2)	3
MS-Register	4
MS-Release	5

Figure 5

OMNI Option (2)

- Type 1/2 ifIndex-tuples include:
  - ifIndex identifies ANET interface
  - ifType is ANET interface type
  - Provider ID is ANET service provider
  - Link quality metric - ‘1’ (lowest) to ‘15’ (highest); ‘0’ means link is DOWN
- Type 1 includes vector of 2-bit preference values. (P00 – P63) correspond to 64 IP DSCPs; P64 and above are additional traffic selectors. Each P[*] encodes (3/2/1/0) for high/medium/low/disabled
- Type 2 includes RFC6088 format traffic selector (flow bindings coded in IPv6 ND messages w/o requiring adjunct mobility messages)

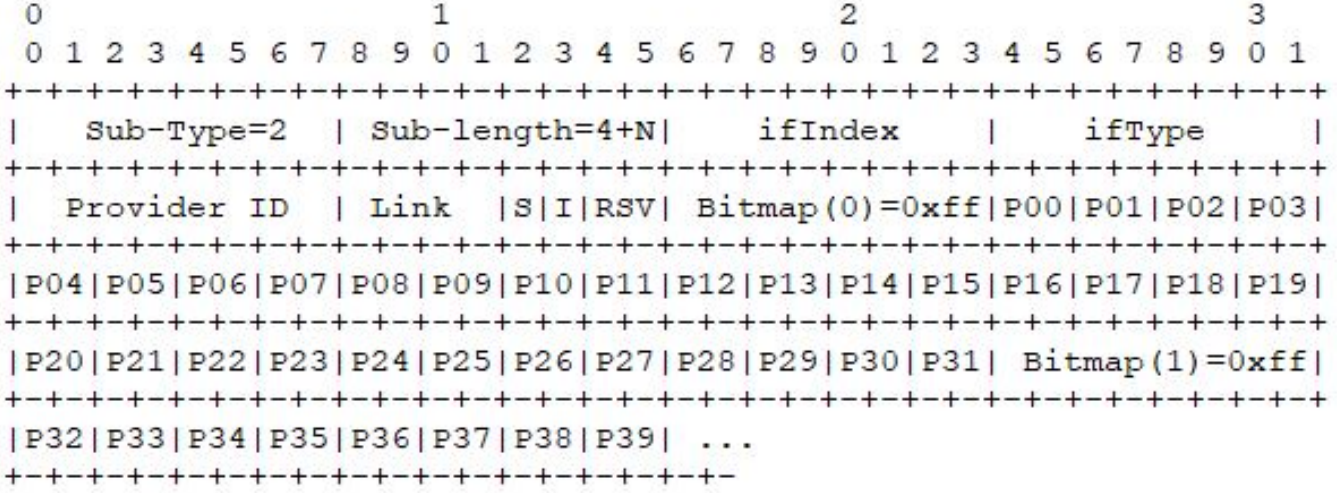

Figure 8: ifIndex-tuple (Type 1)

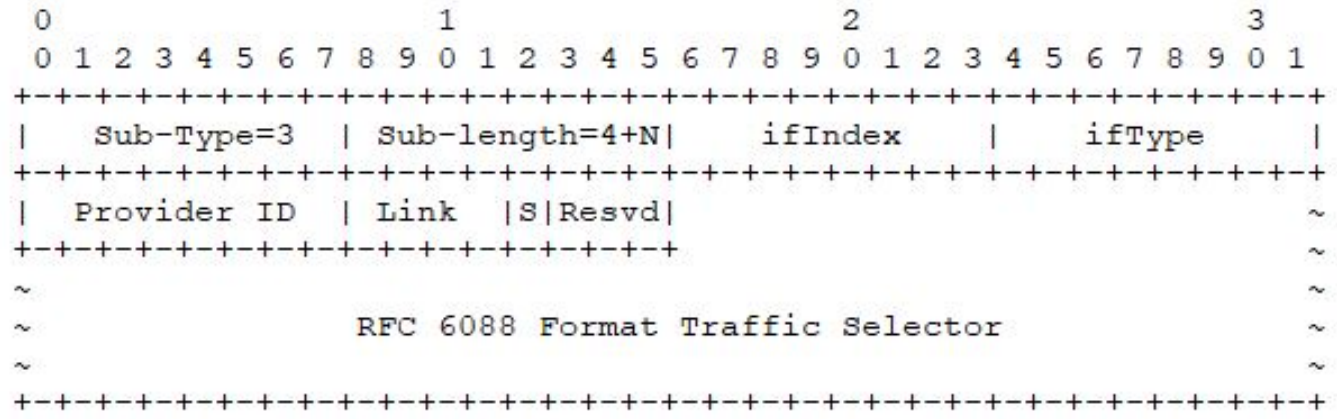

Figure 10: ifIndex-tuple (Type 2)

Router Discovery and Prefix Registration

- MN sends Router Solicitations (RS) with OMNI option
- AR receives RS and coordinates with MS
- AR returns Router Advertisements (RA) with OMNI option; any autoconfiguration information that would normally be delivered in a solicited RA message.
- MN sends initial RS's to register MNP and initial set of UP ANET interfaces:
  - RS messages sent from within the OMNI interface over an UP underlying ANET interface. Process coordinated from within the OMNI interface and opaque to the IPv6 layer.
  - AR processes RS message and conveys the OMNI option information to the MS.
  - MS injects MNP into the routing system then caches Prefix Length, MNP and ifIndex-tuples.
  - MS directs AR to return RA message to MN with an OMNI option and with non-zero Router Lifetime if Prefix Registration successful (otherwise zero Router Lifetime).
  - MN receives RA confirmation; AR will then forward packets between the MN and the MS.

Router Discovery and Prefix Registration (2)

- After initial registration, when ANET interface transitions to UP, MN sends OMNI RS.
- When ANET interface transitions to DOWN, MN sends unsol. Neighbor Advertisement (uNA) over any UP ANET interfaces with ifIndex-tuple for DOWN interface with Link=0.
- When MN departs from a current MSE, sends OMNI RS/uNA over any UP ANET interface with an MS-Release ID.
- When MN associates with a new MSE, sends OMNI RS/uNA over any UP ANET interface with an MS-Register ID.
- When all of a MNs underlying interfaces have transitioned to DOWN (or if prefix registration lifetime expires) MS withdraws the MNP.
- MN responsible for retrying RS until an RA is received. If no RA is received over multiple UP ANET interfaces, MN declares this MSE unreachable and tries a different MSE.

Conceptual Multilink Selection Algorithm

- MN's IPv6 layer selects the outbound OMNI interface according to standard IPv6 conceptual sending algorithm
- OMNI interface maintains default router list and neighbor cache
- For packets originating from the MN, OMNI interface selects downlink ANET interface based on multilink parameters, e.g, DSCP, application port number, cost, performance, message size, etc.
- For packets in the network destined to MN, MS selects an uplink ANET interface based on OMNI option traffic selectors
- Reordering and duplication – honors section 15 [RFC3819]

Status and Next Steps

- Work-in-progress in WG-I Mobility Subgroup since March 2019 – now published as IETF Internet Draft (draft-templin-6man-omni-interface)
- Code in advanced stages of development and testing in linux-based network emulations and early aviation testbed trials (also in other relevant trials, e.g., enterprise network mobile devices)
- Liaison statement requesting IETF to adopt the work:
  - Example of an “IPv6-over-(foo)” spec
  - Uses new IPv6 ND option (the OMNI option)
  - Uses new link-local address format (the OMNI LLA)
  - Standards-Track RFC requested

➢ **Additional review input / co-authors welcome**

➢ **Adopt as 6man Working Group Item?**

Backups

OMNI Interface Benefits

- no need to obtain a “care-of” address on each physical interface, and the same unique MNP-based IP addresses can be used natively and without encapsulation over all physical interfaces
- no need for over-the-air messaging for Duplicate Address Detection (DAD), Multicast Listener Discovery (MLD), Address Resolution and Prefix Delegation and no need for over-the-air encapsulation.
- can amount to a significant savings in avoiding costly and unnecessary radio transmissions.
- provides a nexus for multilink coordination guided only by the information included in the OMNI option in IPv6 ND messages and without need for additional adjunct protocol messaging.

OMNI Interface Benefits (2)

- In the network, using the OMNI virtual overlay model provides for distributed mobility management where the load for coordinating mobility is spread between potentially many mobility service agents instead of concentrated at only one or a few agents.
- A single MN to ground RS/RA message exchange can also be used to coordinate multiple network service endpoints in point-to-multipoint fashion.
- High availability and resilience are naturally supported, as well as true path MTU diversity while spanning multiple aeronautical network service providers for seamless worldwide mobility.