OMNI Adaptation Layer (OAL)

IETF110 6MAN Working Group - March 9, 2021

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

The Boeing Company

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

Overlay Multilink Network Interface (OMNI)

- draft-templin-6man-omni-interface
- Overlay interface configured over multiple underlying interfaces:

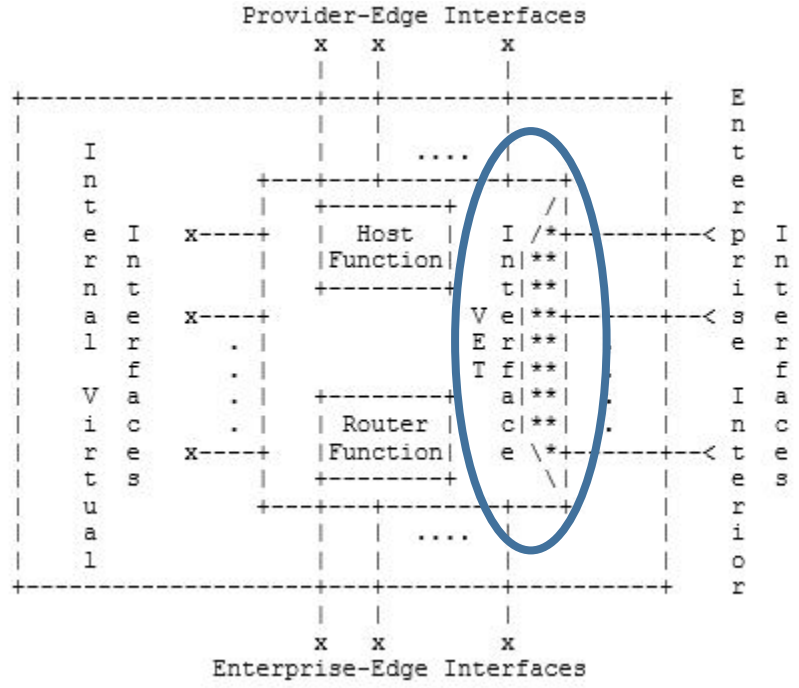

OLD: (from RFC5558, February 2010)

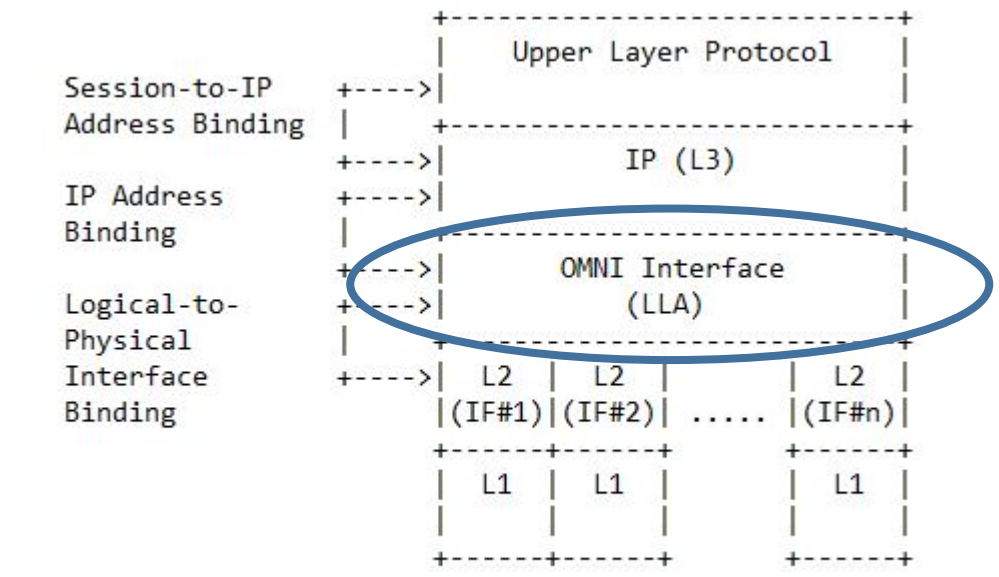

Figure 1: OMNI Interface Architectural Layering Model

NEW: (Copied from RFC7847, May 2016)

OMNI Interface Characteristics

- Ordinary interface with 9180 MTU (IP layer expects interface to deliver packets/fragments up to 9180 bytes)
- Internally performs IP encapsulation to convey original IP packets up to 9180 bytes over diverse underlying interfaces
- Underlying network path MTUs often (much) smaller:
  - IPv6 minimum path MTU 1280 (no network fragmentation)
  - IPv4 minimum path MTU 576 (minimum IPv4 interface MTU is 68, but network can fragment and all IPv4 destinations reassemble at least 576)
- Need to “adapt” OMNI interface MTU to underlying network path MTUs – The OMNI Adaptation Layer (OAL)

OMNI Adaptation Layer (OAL)

- OMNI interface sublayer below IP but above underlying interfaces based on RFC2473 encapsulation
- When IP layer delivers packet to OMNI interface, OAL source inserts RFC2473 encapsulation header and appends 2 byte trailing Fletcher checksum to form “OAL packet” (trailer counted as part of payload):

- OAL source next uses IPv6 fragmentation to break OAL packet into fragments containing no more than Maximum Payload Size (MPS):

OAL Fragmentation and Reassembly

- OAL source encapsulates each OAL fragment in underlying network headers (e.g., UDP/IP) and sends over underlying interface

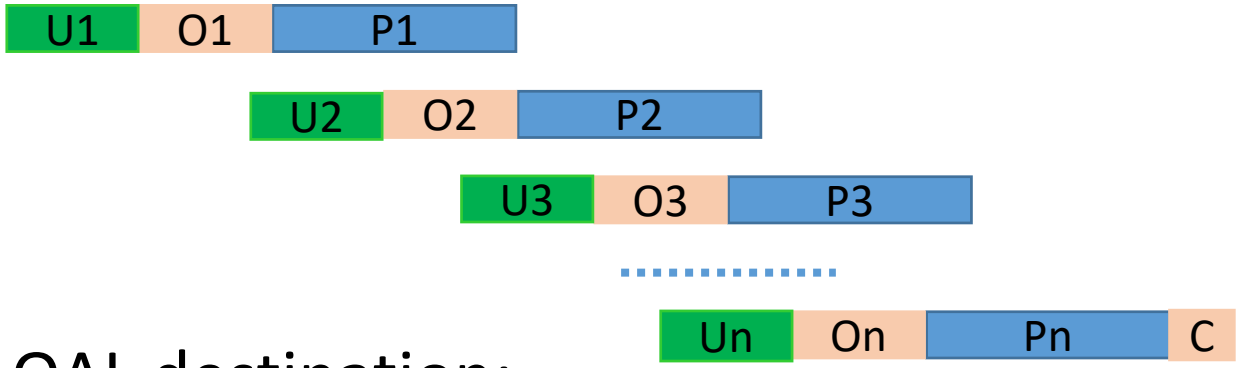

- OAL destination:
  - discards underlying network headers
  - reassembles OAL packet and verifies checksum
  - discards OAL header and trailer and delivers original packet to IP layer

OAL Maximum Payload Size (MPS)

- Some hops in IPv6 OAL destination "path" could be tunnels over IPv4, IPv6/IPv4 translators, etc. – could also traverse concatenated Internetworks with diverse IP protocol versions (more later)
- **IPv4 minimum path MTU (576) assumed unless there is better knowledge**
- OAL encapsulation worst-case (88 bytes):
  - 40 byte RFC2473 header, plus 40 byte OAL Routing Header (ORH) plus 8 byte Fragment Header
- Underlying network encapsulation worst-case (88 bytes):
  - 40 byte IPv6 header (or, 20 byte IPv4 header), plus 40 bytes for security encapsulations (IPsec, SSL/TLS, etc.), plus 8 byte UDP header
- **Minimum MPS (minMPS) therefore (576 – 88 – 88) = 400 bytes**
- Example: worst-case for 1500 byte original IP packet is 4 OAL fragments
  - 3x fragments with 400 byte payloads
  - 1x fragments with 302 byte payload (includes 2 octet trailer)
- **Fortunately, larger per-path MPS values can often be determined**

OAL Maximum Payload Size (MPS) (2)

- OAL sources can set “path MPS” values larger than minMPS for specific OAL destinations
- If OAL source knows (i.e., without probing) that path can transit larger MPS without loss, it can set a larger value for that OAL destination
- OAL source can send probes to OAL destination to discover larger path MPS (RFC4821/RFC8899)
- OAL encapsulation not needed when source and destination are on the same link and original IP packet fits within the link MTU

OAL Packet/Fragment Validation

- minMPS safe assumption that works over all paths - non-final OAL fragments must contain at least minMPS worth of payload
- OAL extension headers may include one Fragment Header and one ORH, but no other IPv6 extensions headers
- OAL destinations drop all non-final OAL fragments with less than minMPS of payload - defeats tiny fragment attacks
- OAL destinations drop all OAL packets/fragments with OAL extension headers other than a single Fragment Header and a single ORH

OAL Addressing

- RFC2473 header requires IPv6 source/destination addresses
- RFC4193 ULA-D's used as OAL source/destination to enable forwarding at a layer below IP.
- From perspective of inner IP layer, OAL forwarding indistinguishable from Layer-2 bridging
- Can be used to traverse multiple independent Internetworks "concatenated" by bridges

OAL Single Network Traversal

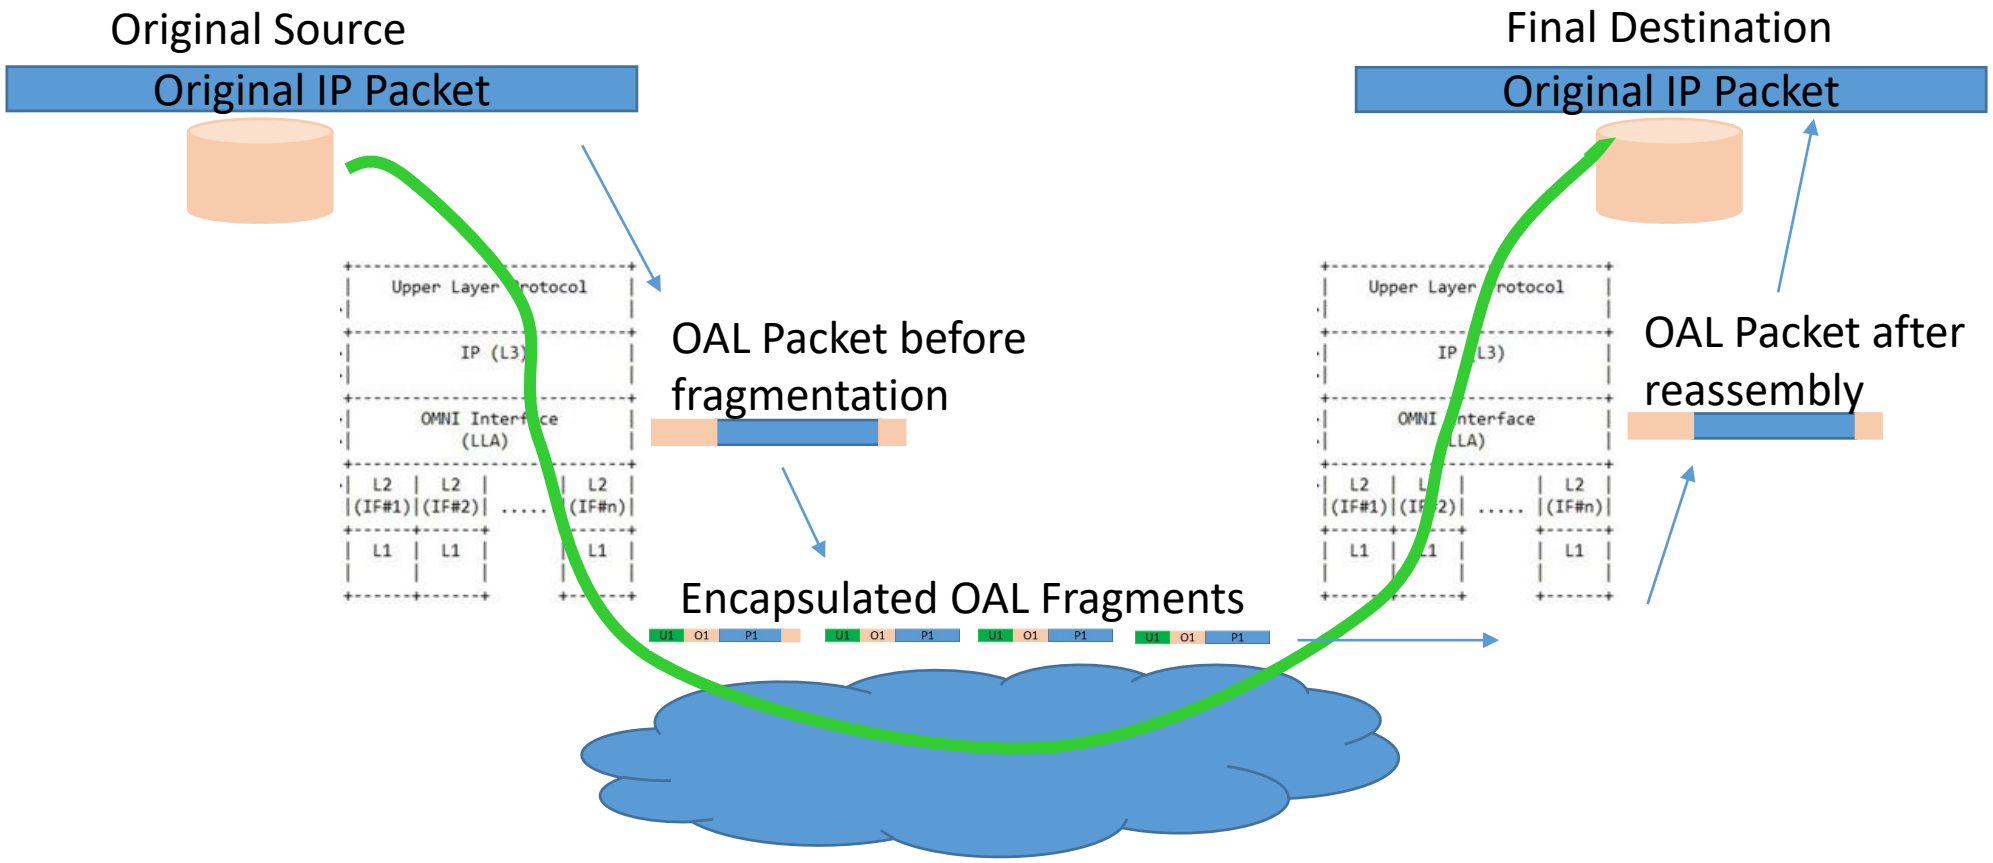

OAL Multi-Network Traversal

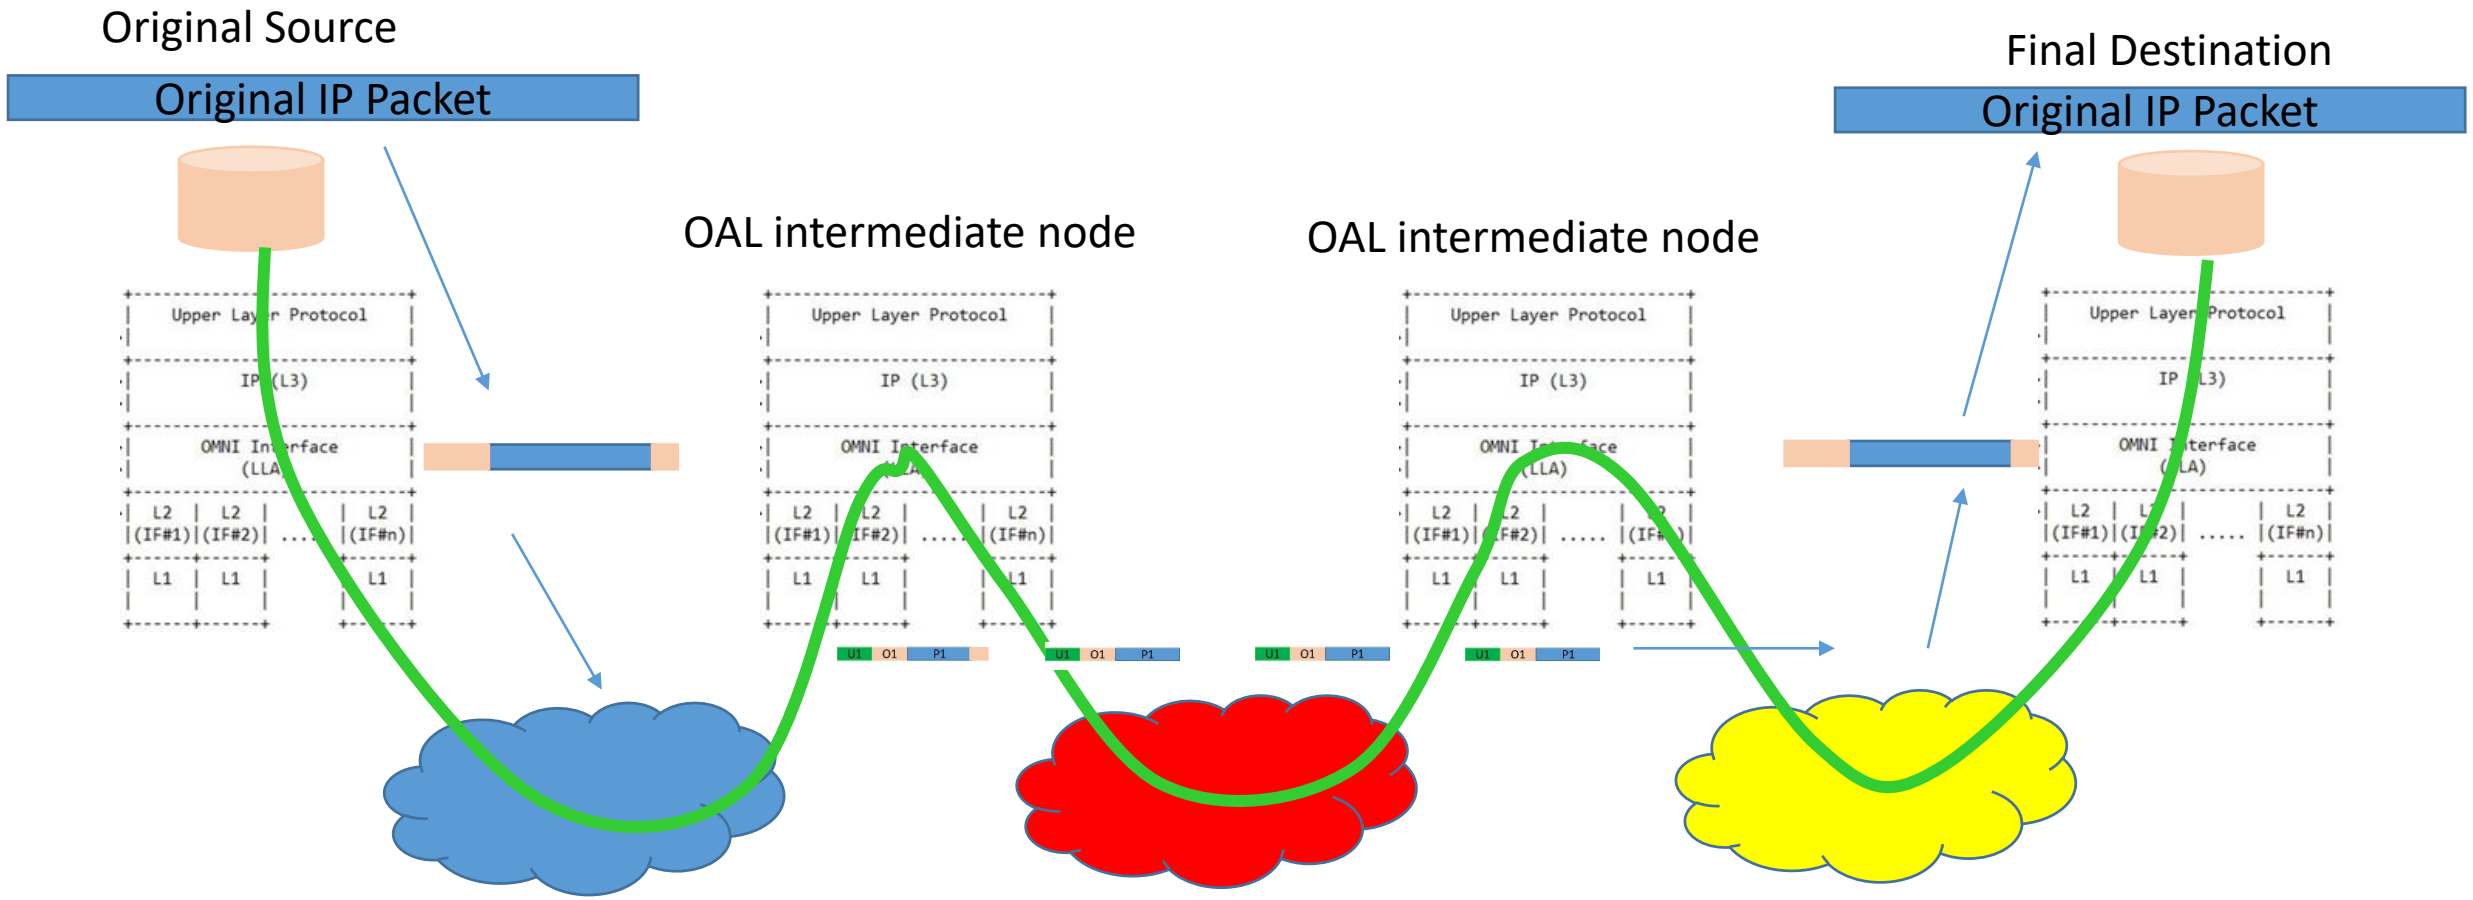

OAL Super Packets

- Original IP packets often smaller than the OMNI interface MTU and arrive in “bursts”
- Original IP packet transmission sometimes triggers control message to same OAL dest.
- May be more efficient to “pack” multiple original IP packets; control messages into a single OAL “Super Packet”

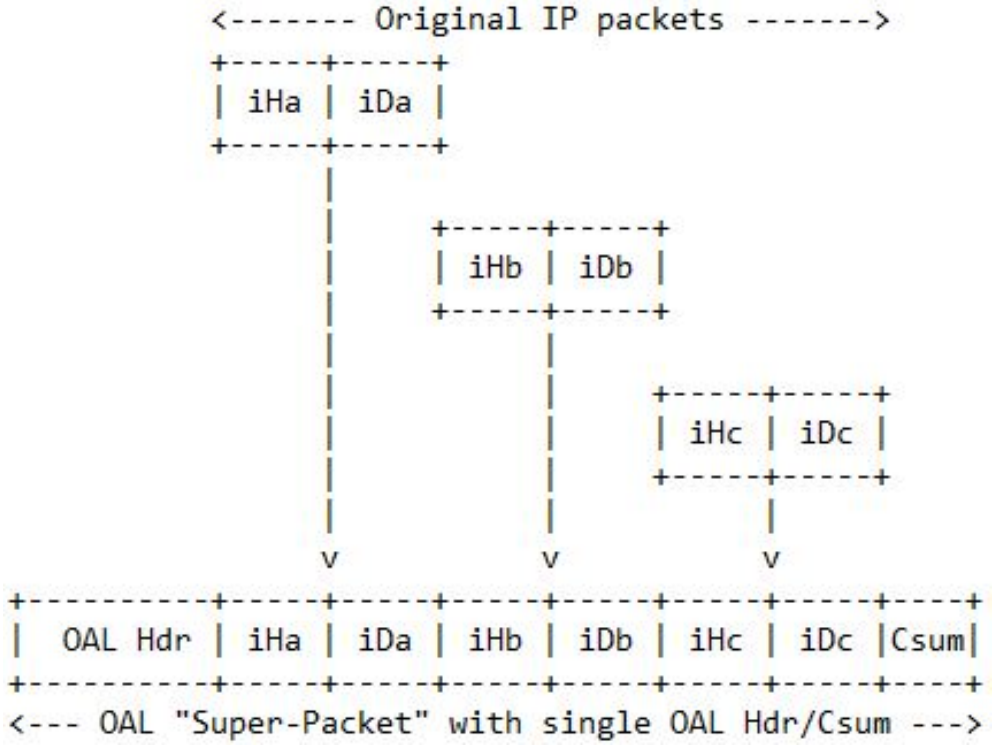

(Copied from ‘draft-ietf-intarea-tunnels’)

OAL Packet Size Feedback

- Although OMNI interface accommodates packets up to 9180, not always good to continuously send such large packets
- Classic PMTUD sends Packet Too Big (PTB) **“hard errors”** to inform sources of packet loss due to size restrictions
- OMNI interface continuously forwards packets up to 9180 while sending PTB **“soft errors”** – results in **Lossless PMTUD**
- New capability for hosts to dynamically tune packet sizes for optimal performance without loss

OAL Integrity

- OAL is a new sublayer, hence must include its own integrity check
- Underlying network hops use CRC-32 and upper layers use Internet Checksum – OAL uses Fletcher for multi-layer diversity
- Fletcher also good for detecting reassembly misassociations, which are critical for encapsulations over IP that may incur fragmentation
- Underlying network UDP checksums can be disabled for OAL-encapsulated packets, but still needed for non-OAL packets and control messages when address/port integrity is required [RFC6935][RFC6936].
- Some underlying network hops (e.g., tunnels over IPv4) may not include integrity checks. OAL checksum detects unprotected underlying hop corruption; improves Internet integrity over current state of affairs.

Backups

OAL Bridging of Multiple Network Segments

- OMNI “Link” Consists of “Segments” joined by OAL Intermediate Nodes acting as “Bridges”
- Example: Civil Aviation has multiple providers including ARINC, SITA, Inmarsat, others

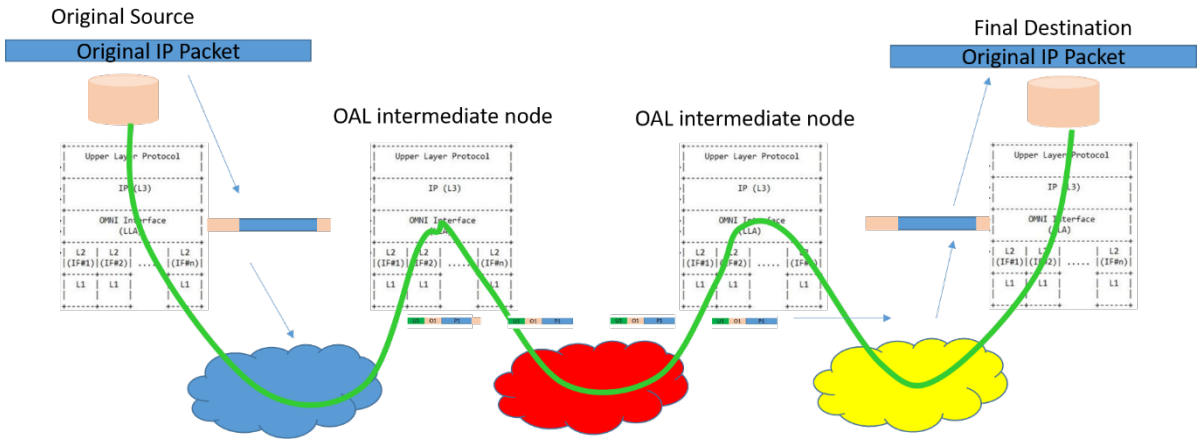

- Second example: bridging network segments within an enterprise network
- Third example: bridging multiple enterprises (Boeing, Airbus, Lockheed, etc.)
- An even more relevant example for this group:
  - **Bridging the IPv4 and IPv6 Internets**