Node Liveness Protocol

Tony Li

Problem

- People have mistakenly used LSP presence as a substitute for node liveness.
- This doesn’t work when the network graduates to hierarchy.
- Proposals on the table:
  - Prefix Unreachable Announcement (PUA)
  - Put loopbacks in BGP
  - Don’t aggregate loopbacks
  - Many more…
- Root problem: IGPs carry reachability, not liveness. Need another mechanism. The IGP is not a dump truck. It doesn’t carry everything.


JOHN DEERE
DEERE
310E

Node Liveness Protocol

- Publish-subscribe service
- Runs on ABRs
- Simple register/notification messaging
- Could support much more than liveness (e.g., non-IGP router capabilities)

Usage

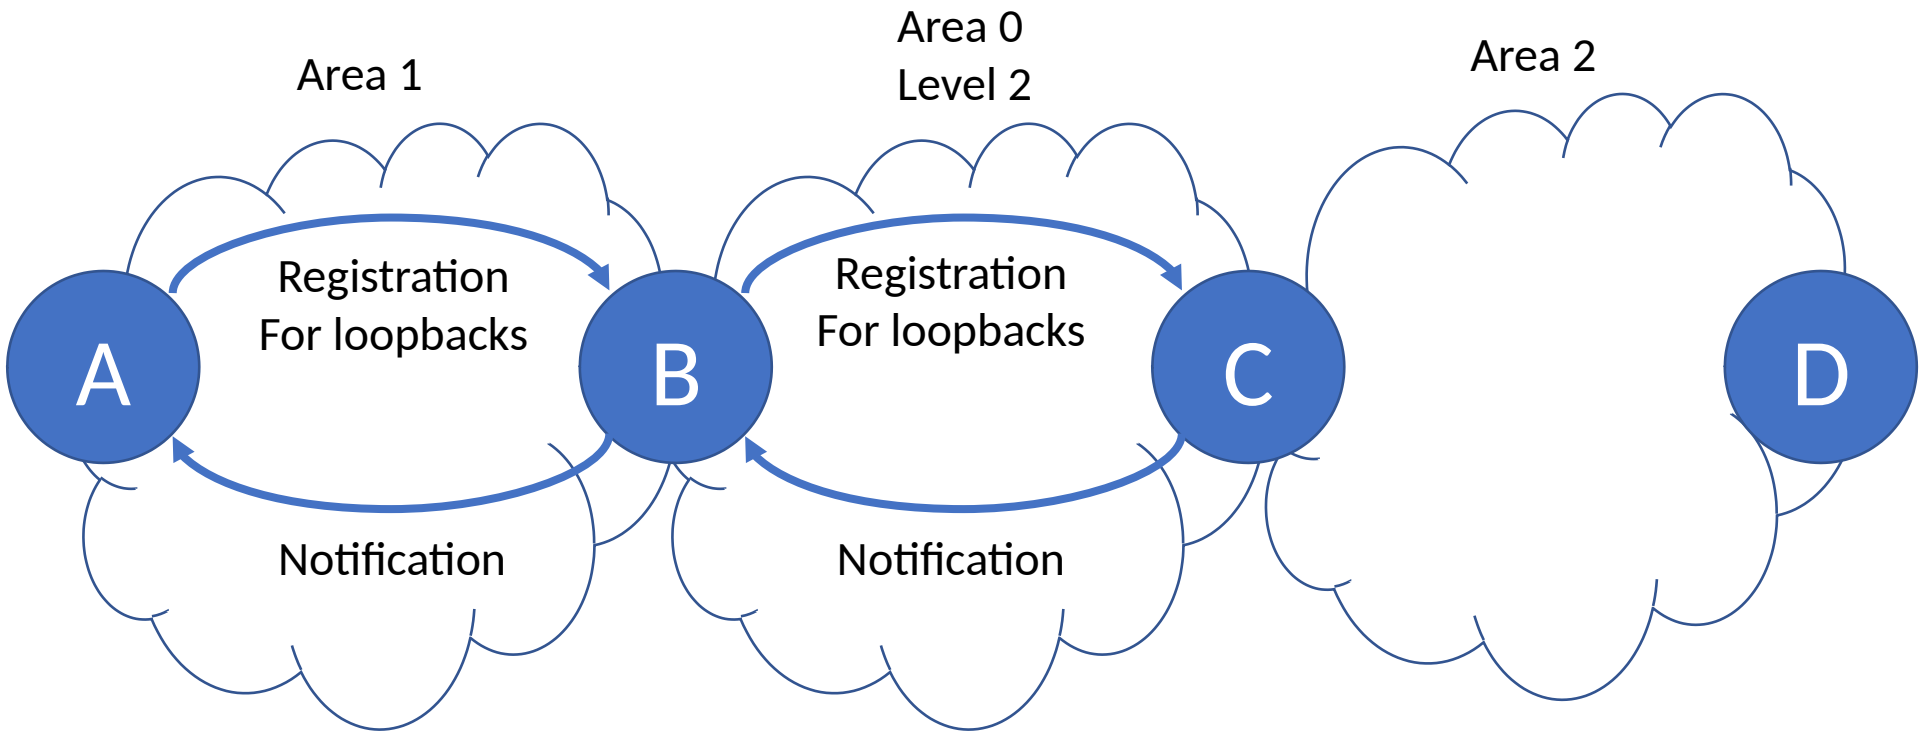

Details

- Put all loopbacks under a single, hierarchical prefix
- Local ABR registers for all remote more-specific prefixes
- With N nodes per area and A areas, 2 ABRs/area, an ABR sees:
  - N local sessions and registrations
  - 2(A – 1) remote sessions and registrations
- N=1000, A=1000: 1000 + 2 * 999 = 2998
- If you don’t want to run this on an ABR, any IGP speaker in the area will work; that’s the ‘dump truck’
- Added IGP capabilities to identify NLP speakers.

Registration message

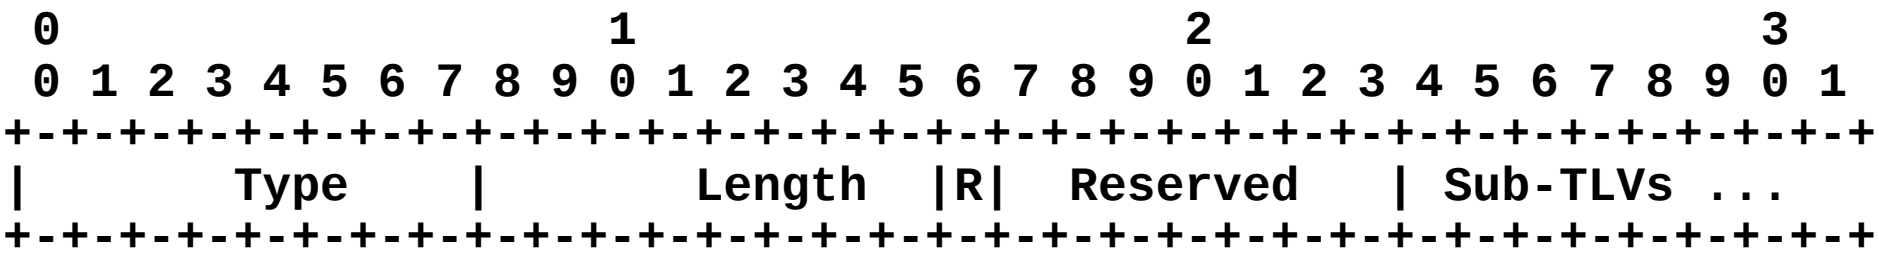

Type: 1 (Registration Message), 1 octet

Length: 1 + length of the sub-TLVs, 1 octet

R: 1 bit

- **0: Register**
- **1: Unregister**

Liveness registration sub-TLV

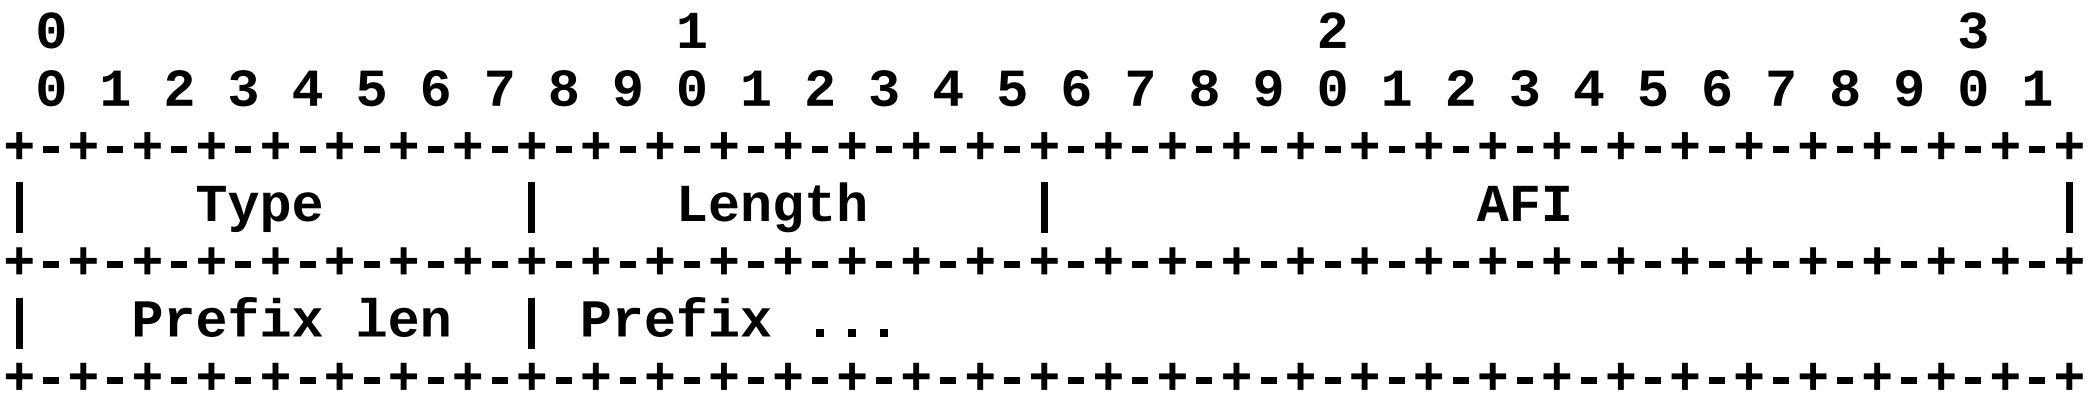

Type: 1 (Liveness), 1 octet
Length: 3 + the number of octets for the prefix, 1 octet
AFI: Address Family Identifier [afireg], 2 octets
Prefix len: number of significant bits in the prefix, 1 octet
Prefix: The prefix to register/unregister, n octets

Notification Message

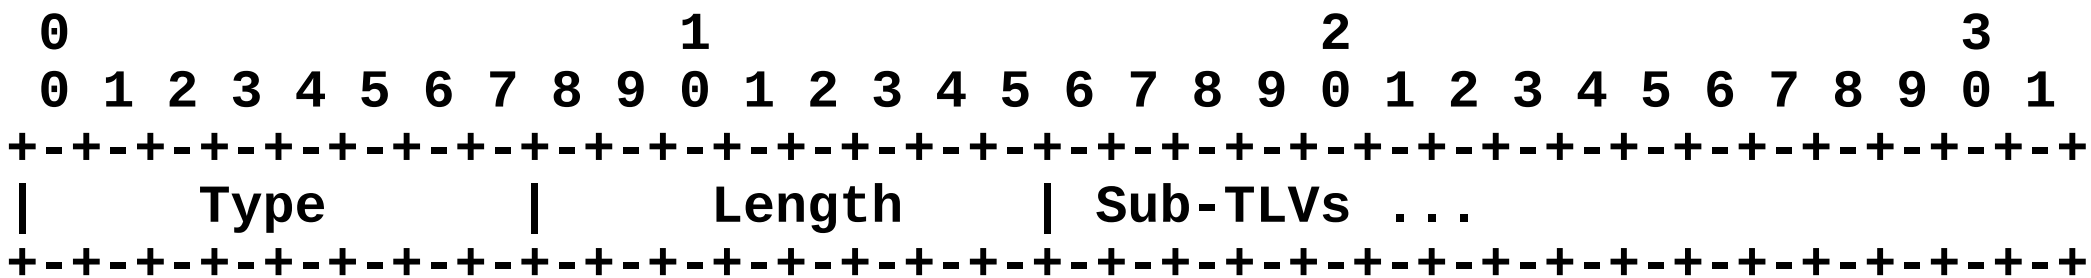

Type: 2 (Notification), 1 octet
Length: length of the sub-TLVs, 1 octet
Sub-TLVs: One or more sub-TLVs, specifying the registration/ unregistration. Variable length.

Liveness notification

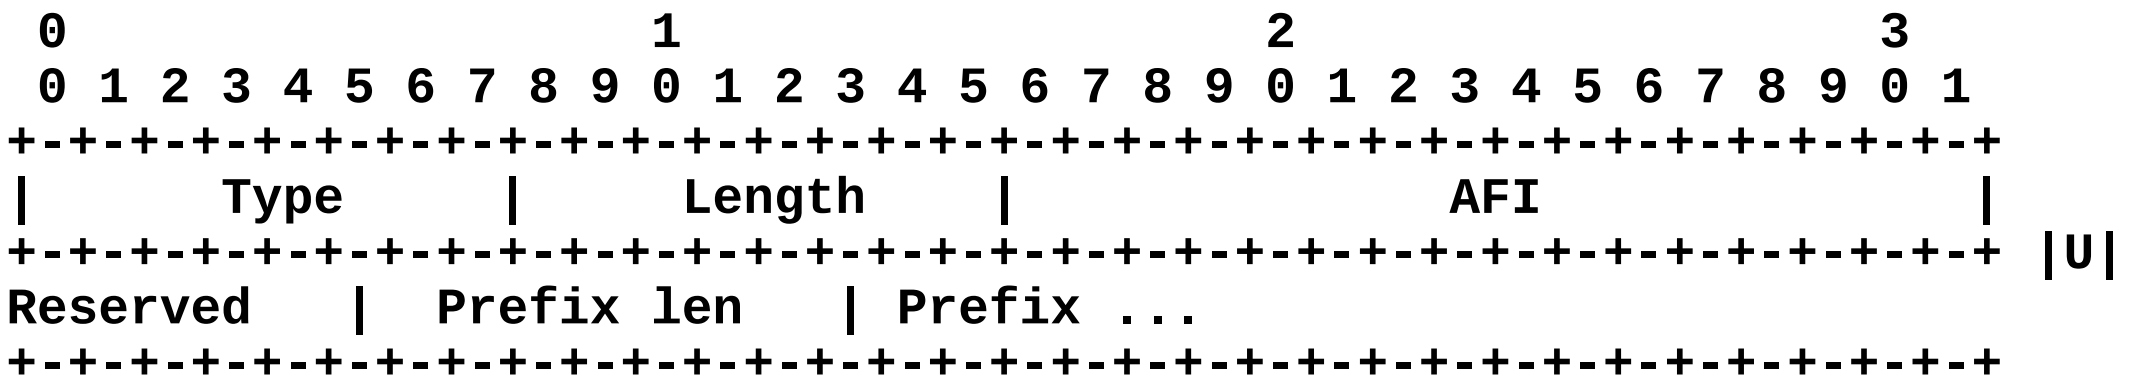

Type: 1 (Liveness), 1 octet
Length: 3 + number of octets of prefix, 1 octet
AFI: Address Family Identifier [afireg], 2 octets
U: 1 bit 0: Up Event 1: Down Event
Reserved: must be zero and ignored on receipt, 7 bits
Prefix len: number of significant bits in the prefix, 1 octet
Prefix: The prefix generating the notification, n octets

IS-IS capability

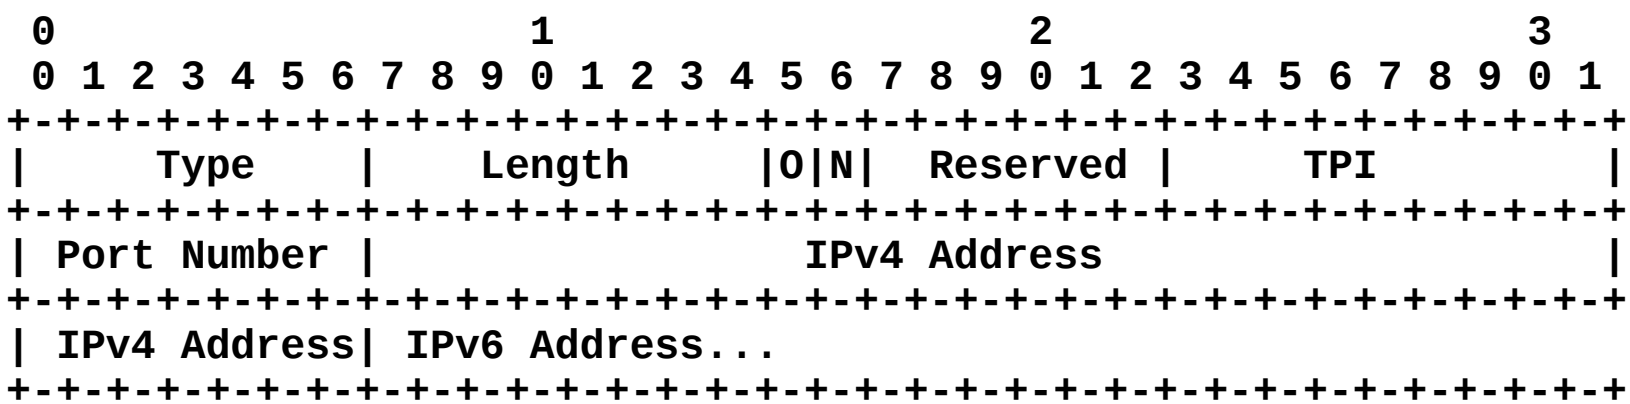

Type: TBD1
Length: n * (4 octets + 4 octets if O is set + 16 octets if N is set) O: 1 if an IPv4 Address is included
N: 1 if an IPv6 Address is included
Reserved: must be zero and ignored on receipt, 6 bits
TPI: Transport Protocol Identifier, 1 octet 0: TCP 1: QUIC
Port Number: Transport protocol port number, 2 octets
IPv4 Address: Service contact address, 4 octets if the O bit is set, 0 otherwise.
IPv6 Address: Service contact address, 16 octets if the N bit is set, 0 otherwise.

OSPF capability

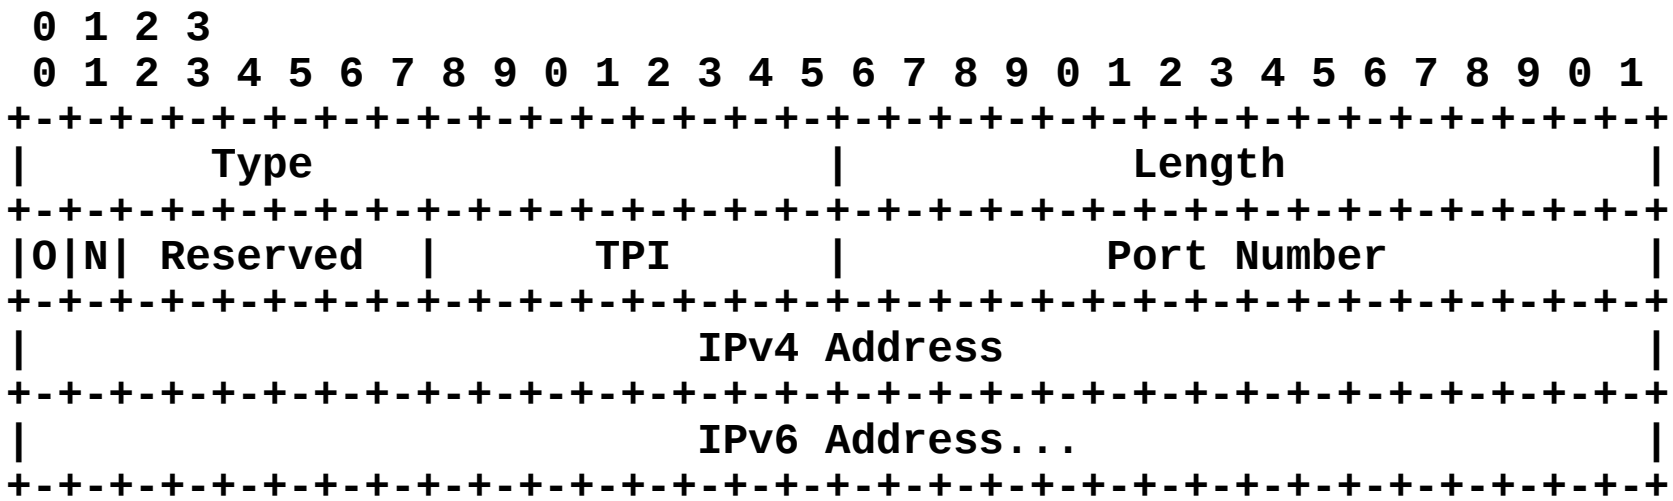

Type: TBD2
Length: n * 3 octets
O: 1 if an IPv4 Address is included
N: 1 if an IPv6 Address is included
Reserved: must be zero and ignored on receipt, 6 bits
TPI: Transport Protocol Identifier, 1 octet 0: TCP 1: QUIC
Port Number: Transport protocol port number, 2 octets
IPv4 Address: Service contact address, 4 octets if the O bit is set, 0 otherwise.
IPv6 Address: Service contact address, 16 octets if the N bit is set, 0 otherwise.