Multi-MTU subnets

draft-van-beijnum-multi-mtu-05

6man @ IETF-95

Previous presentations

- https://www.ietf.org/proceedings/69/slides/intarea-6.pdf
- https://www.ietf.org/proceedings/71/slides/intarea-4.pdf
- https://www.ietf.org/proceedings/78/slides/intarea-5.pdf

Jumboframes

- Lots of gigabit ethernet equipment and hosts support larger packets: "jumboframes"
- Common value: ± 9000 bytes
  - but no standard jumboframe size (but RFC1626 (SMDS) and 2225 (IP over ATM) defines IP MTU 9180)
- "Mini jumbos" / "baby giants" up to ± 2000 bytes common in lower-speed switches

Subnet with mixed MTUs

- Shared L2 between wireless and wired nodes in residential deployments.
- Wireless PHY generally only support ~2000 MTU.
- Wired PHY generally supports ~9000.
- Default is 1500.

What we need

- Ability to turn on jumbos without touching all hosts on a subnet
- Take advantage of hardware improvements without protocol work
  - no more hardcoding of MTU sizes
- Mechanism for nodes with different MTU to co-exist on the same L2 segment
- Be backward compatible!
  - also with current jumbo deployments

How?

- ND option to get hint of node MTU size
- UDP (or perhaps ICMP) probing protocol to:
  - see if that packet size works
  - if not, probe for a packet size that works
- Monitor sending/receiving of large packets
  - (similar to IPv6 neighbor unreachability detection or Shim6 REAP)

ND NODEMTU

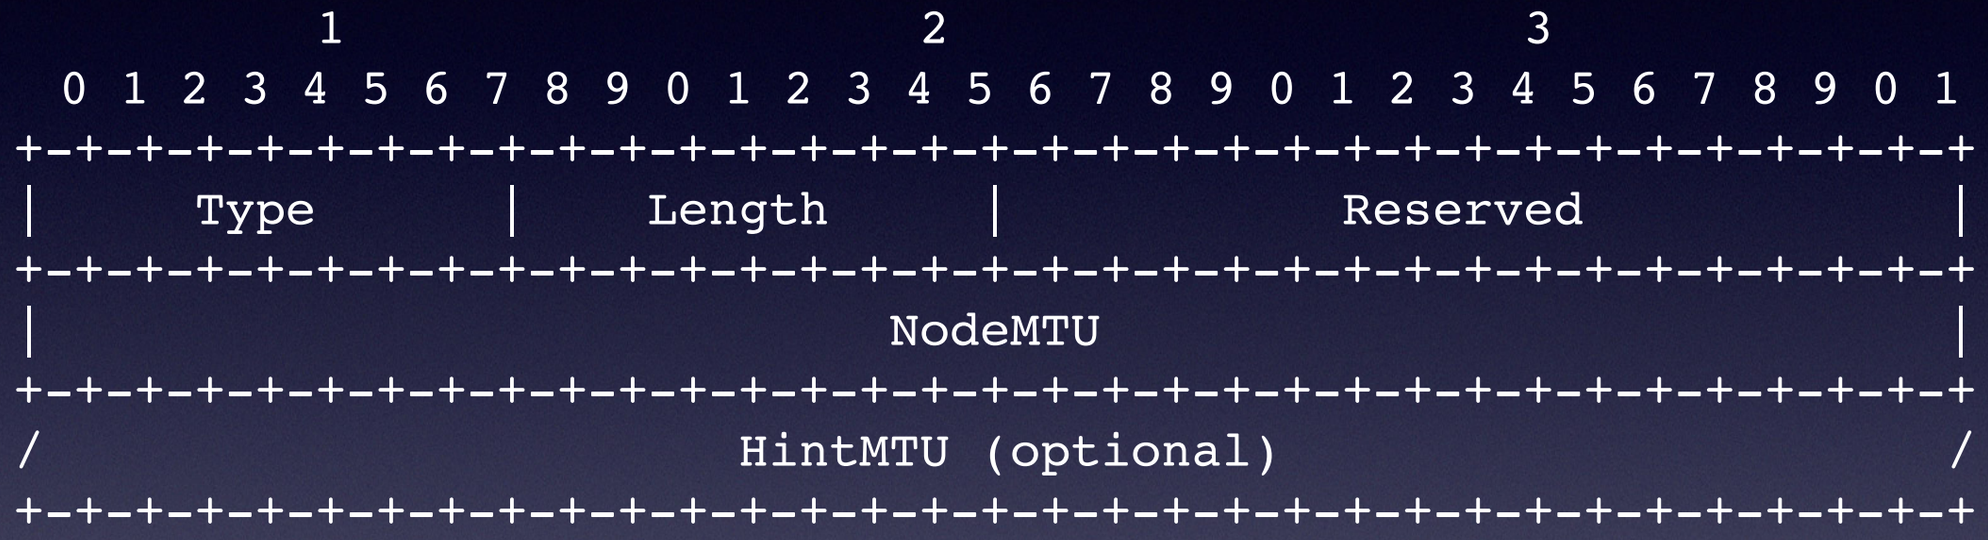
1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
Type
Length
Reserved
NodeMTU
HintMTU (optional)

MTUTEST UDP packet

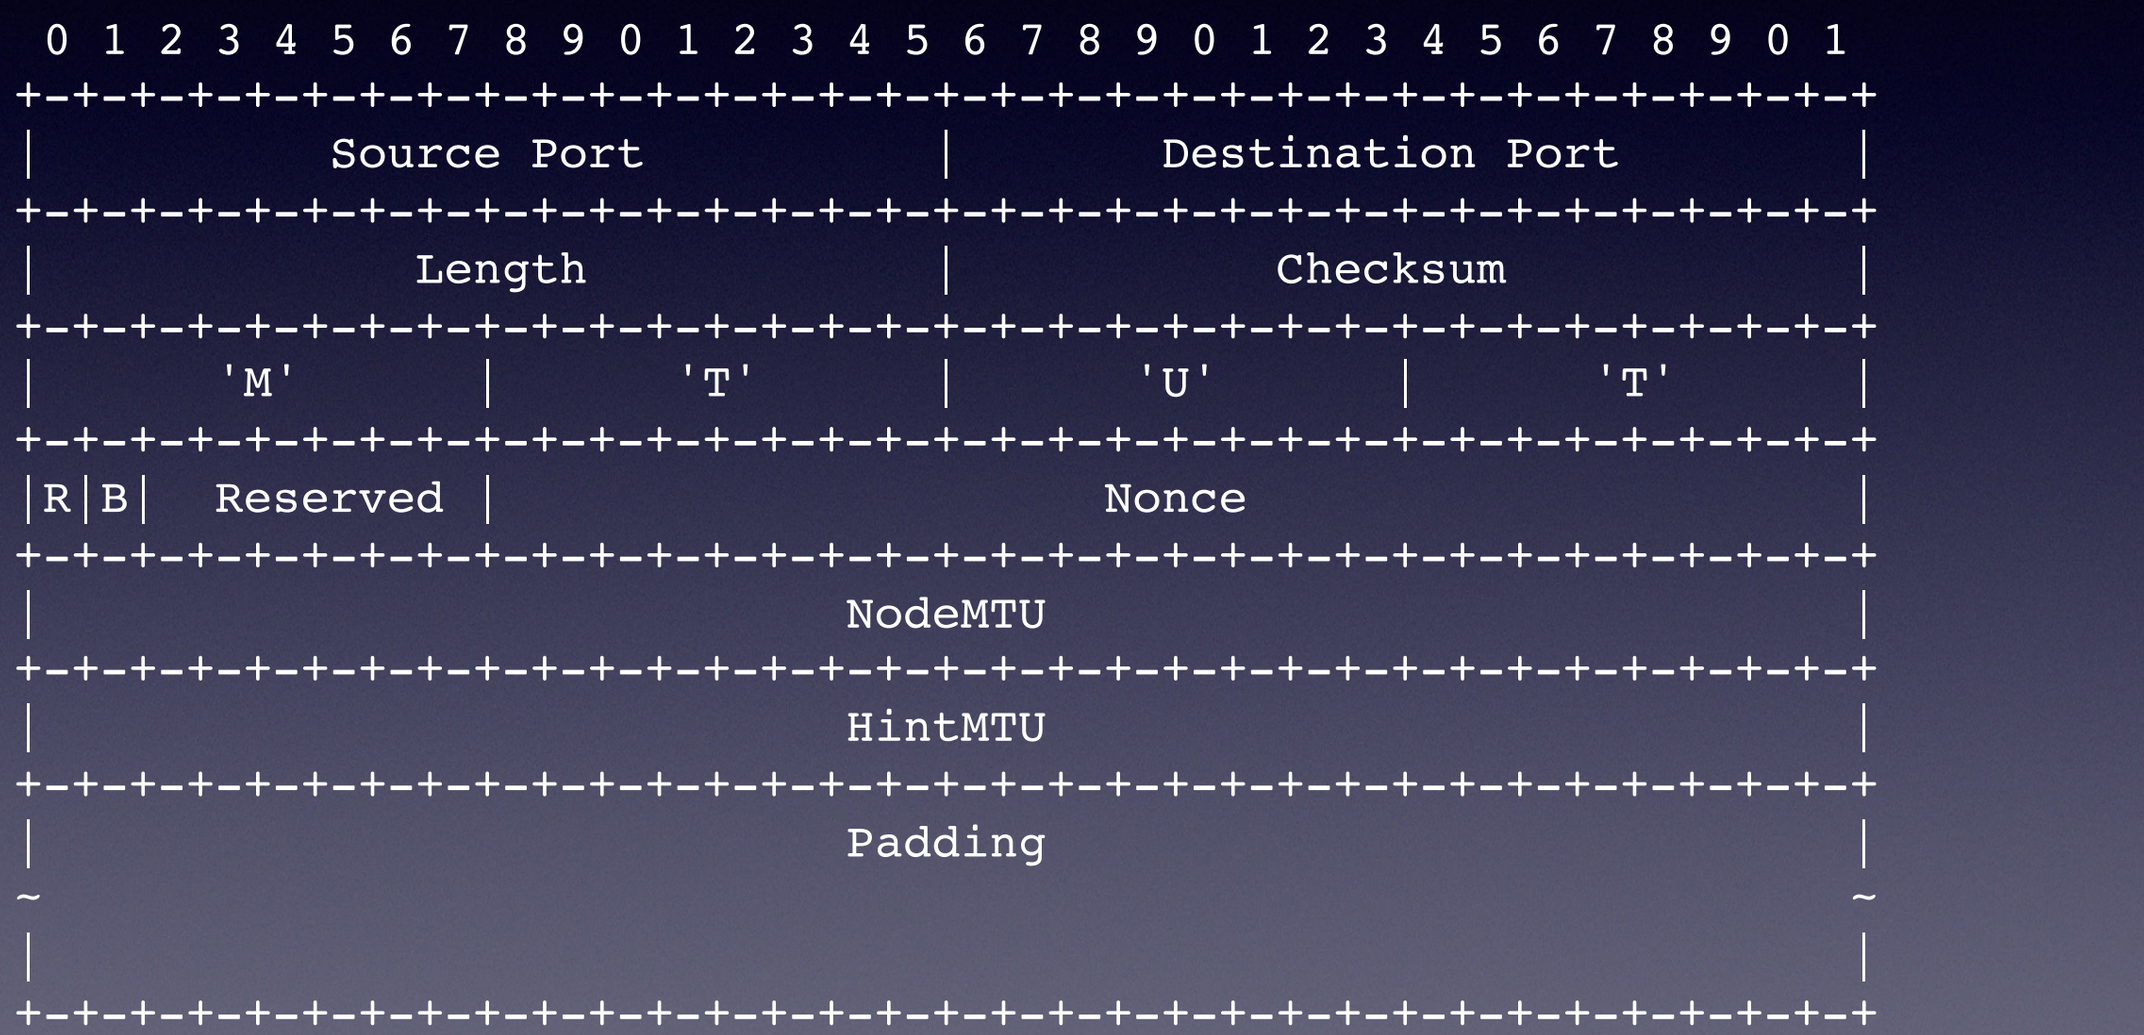

Probing

- Discover capability/remote MTU with minimum size probe
- Establish upper bound quickly:
  - 320, 640, 1280, 2560, 5240, 10240, ...
- Then use hints:
  - 576, 1492, 1500, 1530, 1982, 2304, 4070, 8092, 9000, 16384, 32000, 64000

Thanks, all.

Questions?