



Tunnel Issues ID - Status

Joe Touch, USC/ISI
Mark Townsley, Cisco



Motivation

- Tunnel use common
 - tunnel+MTU+ICMP in ~100 RFCs
 - IPsec, L2TP/PPTP
 - Mobile IP
 - L[1,2,2.5,3,3.5]VPNs
 - SEAL, LISP
- Potential need for automation
 - 1300-byte MTU vs. can/should we do better
- Potential need to revise/coordinate
 - Fragmentation handling, ICMP handling
- **GOAL: explain in a single document**

Observations

- Tunnels are L2
 - We create them
 - Still subject to link issues, e.g., MTU discovery, signalling
- Advantages vs. other L2s
 - Arguably easier to change
 - When L2 protocol matches L3, it MAY be easier to align L2 and L3 MTU discovery, signalling, etc.

Known Issues

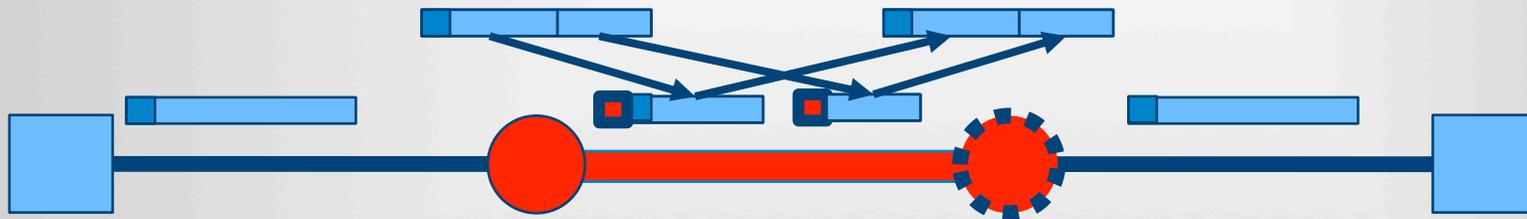
- MTU issues
 - MTU discovery
 - Fragmentation – outer or inner
- Other signalling
 - ICMP
- Performance issues
 - IP-ID exhaustion
 - Fragment size
 - Packing (ala GigE packet bursting)

MTU Discovery

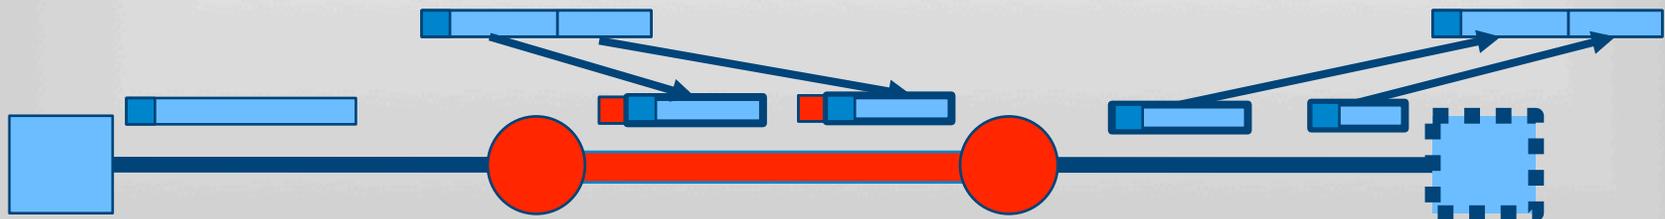
- Mechanisms
 - ICMP-based (RFC 1191)
 - Probe-based (RFC 4821, SEAL)
- Impact on E2E MTU discovery
 - Forwarding/recomputing/validating ICMPs
 - Encapsulator sending advisory too-bigs
- Tunnel MTU discovery
 - Is internal mechanism required?
- See RFC 4459...

Fragmentation

- ***Outer*** implies reassembly at decapsulator

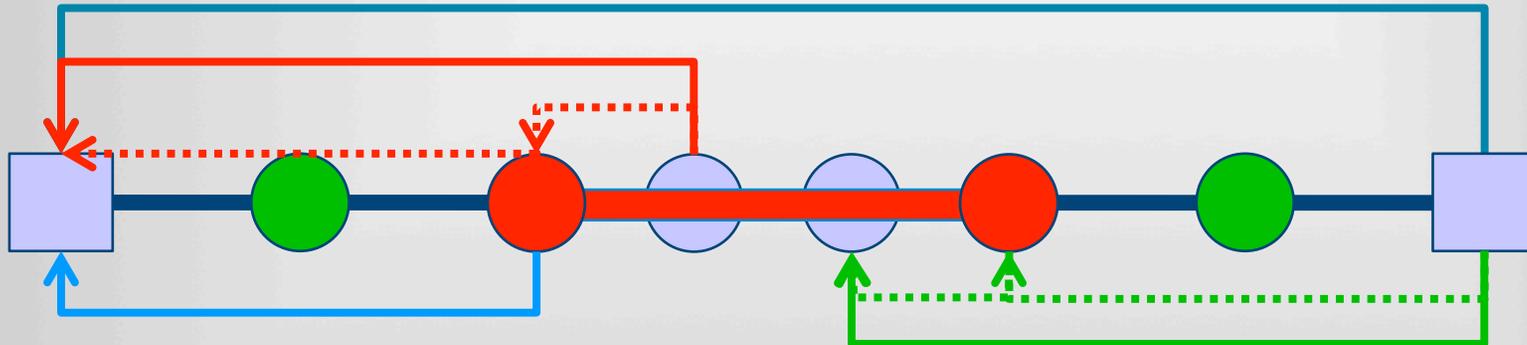


- ***Inner*** affects IPv4 DF, reassy at dst



Signaling – ICMP, etc.

- Pop control out of tunnel?
 - E.g., ICMP underliverables, MTU discovery



- Send tunnel status to the original src?
- Push control into tunnel (ever)?
 - (listed for completeness)

Current Status

- Need contributors
 - Expanded list of examples
 - Placeholder for multipoint
 - Entire section of additional issues
- Relationship to security concerns doc
 - Currently proceeding separately (cross-ref)