Propagating ECN across IP tunnel Headers Separated by a Shim

draft-ietf-tsvwg-rfc6040update-shim-04

Bob Briscoe
bobbriscoe.net
IETF-99, Jul 2017

Recap

Problem (#1) unique to ECN

- Both Diffserv (traffic class) and ECN have to propagate across layers
  - DS propagates 'requirements' down
  - ECN propagates...
    - ECN field down (copy)
    - congestion experienced (CE) up
- forwarded ECN constructed from inner and outer on decap [RFC6040]
- If ECN decap behaviour absent, encap MUST zero ECN outer

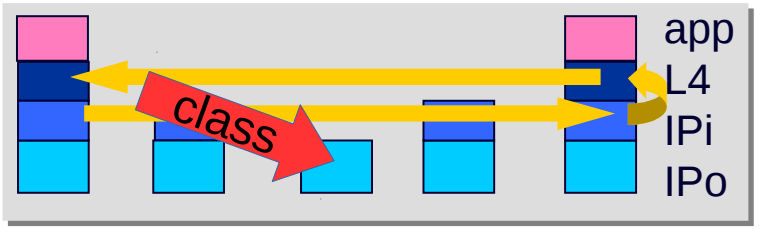

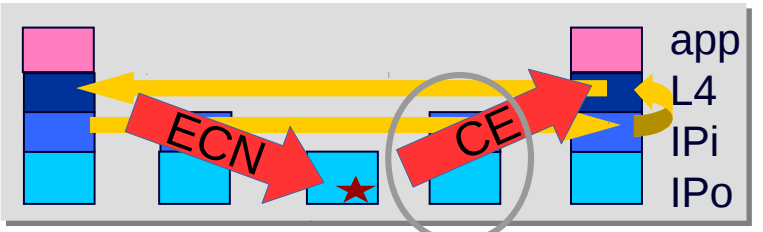

incoming inner	incoming outer			
	Not-ECT	ECT(0)	ECT(1)	CE
Not-ECT	Not-ECT	Not-ECT	Not-ECT	drop
ECT(0)	ECT(0)	ECT(0)	ECT(1)	CE
ECT(1)	ECT(1)	ECT(1)	ECT(1)	CE
CE	CE	CE	CE	CE
	Outgoing header			

✓

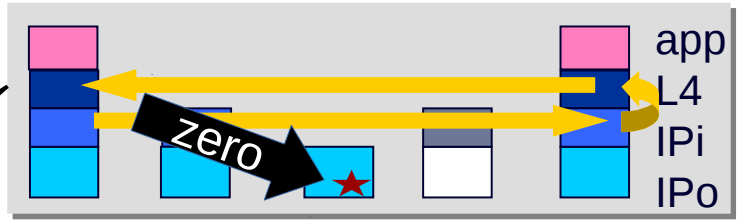

✗

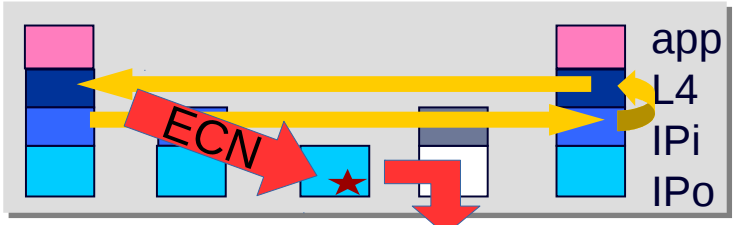

- Non-RFC6040 implementations still have to follow this MUST
  (e.g. by operator config)
- Implementations MUST decouple ECN & DSCP config

draft-ietf-tsvwg-rfc6040update-shim-04

Problem #2

- RFC6040 “Tunnelling of ECN”; scope was “all IP-in-IP tunnels”
- 6040update-shim clarifies
  - scope of RFC6040 includes cases with shim(s)
  - most feasible to propagate ECN with 'tightly coupled shim' (added in same procedure as IP outer)

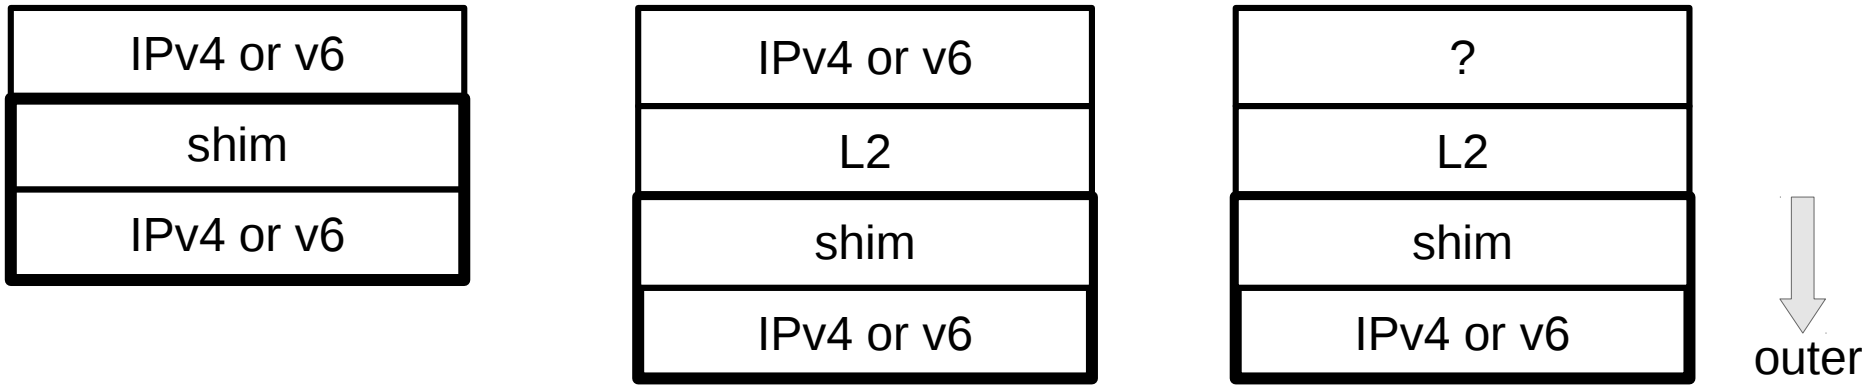

“IP - shim - (L2) - ?IP”

- 6040update-shim is standards track, so it can update standards track RFC6040 and shim tunnel RFCs

Survey of IP-shim-(L2)-?IP encaps

Protocol	RFC	STDs or widely deployed	AOK	NOK: 6040shim updates	NOK: non-IETF: update recommended
Geneve	nvo3-geneve	✓	✓		
GUE	intarea-gue	✓	✓		
UDP Tunnel	8085	✓	✓	✓	
SFC	7665	✓	N/A?		
VXLAN	7348	✓			✓
VXLAN-GPE	nvo3-vxlan-gpe	✗			
LISP	6830	✓	✓		
CAPWAP	5415	✓	✓		
Teredo	4380	✓		✓	
GTP	v1, v1U,v2C	✓			✓
GRE	2784	✓		✓	
NVGRE	7637	✓		⇧	
L2TPv3	3931	✓		✓	
L2TPv2	2661	✓		✓	
PPTP	2637	✗			
AYIYA	www.sixxs.net	✗			
6a44	6751	✗			
SEAL	5320	✗			

Status and Next Steps
draft-ietf-tsvwg-rfc6040update-shim-04

- 4 revs in last IETF cycle
- Thank you Tom Herbert, Joe Touch, Mohamed Boucadair, Carlos Pignataro and Ignacio Goyret, Praveen Balasubramanian
- Milestone: tsvwg WGLC Oct 2017
  - joint with int-area and l2tpext. Any other WG?
- Been pushing to meet that, still feasible

- Open issues

  1) Is SFC really “not applicable”?
  2) Is it true that there are no automated GRE tunnel set-up protocols?
  3) Teredo open issue: mtg next week to close off

Propagating ECN across IP tunnel Headers Separated by a Shim

draft-ietf-tsvwg-rfc6040update-shim-04

Q&A

Updates text for standards track tunnel RFCs

- General ACKs: Alia Atlas for helping to widen then narrow the list
  Tom Herbert, Joe Touch and Mohamed Boucadair
- L2TPv2 & L2TPv3
  - discussed at length on l2tpext list
  - ACK: Carlos Pignataro and Ignacio Goyret
  - written update text to refer to RFC 6040
  - defined and written IANA registry text for L2TP attribute-value-pair (AVP) for tunnel initiator to agree ECN capability with remote tunnel endpoint
- GRE
  - update text refers to RFC 6040
  - no response to questions on int-area list
  - “is it true that there are no automated GRE tunnel set-up protocols?”
- Teredo
  - update text refers to RFC 6040
  - ACK: Praveen Balasubramanian (Christian Huitema was original author, but just left company)
  - open question on tunnel setup – resolution in progress