

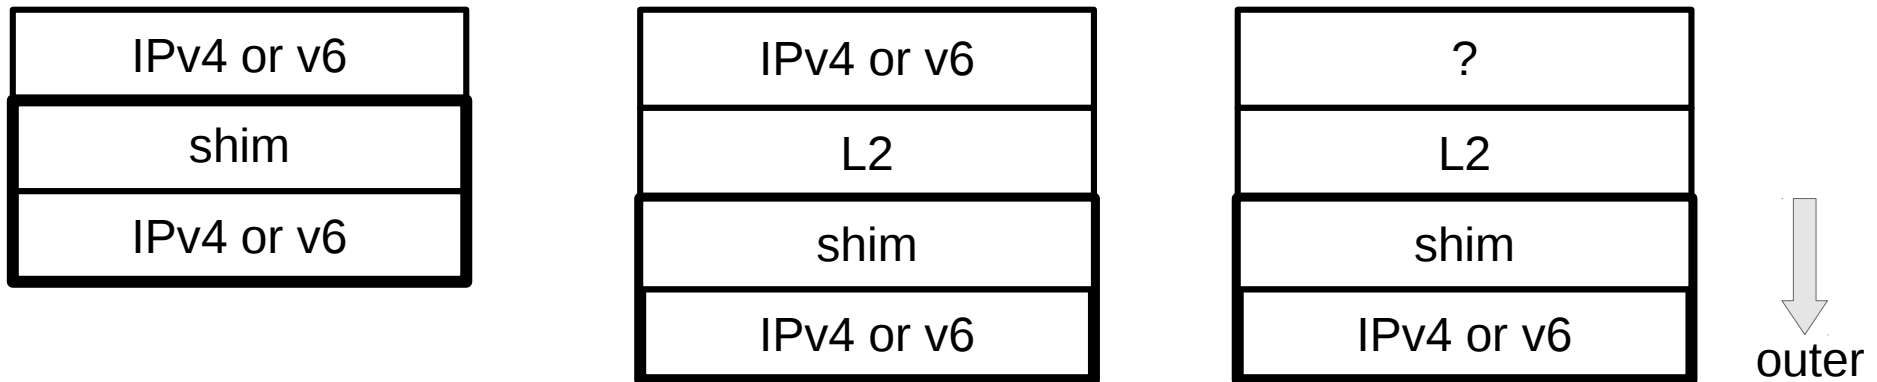
Propagating ECN across IP tunnel Headers Separated by a Shim

draft-ietf-tsvwg-rfc6040update-shim-04

Bob Briscoe
bobbriscoe.net
IETF-99, Jul 2017

Problem #1

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



- 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	✓	✓		
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	✗			

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



Problem #2: 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



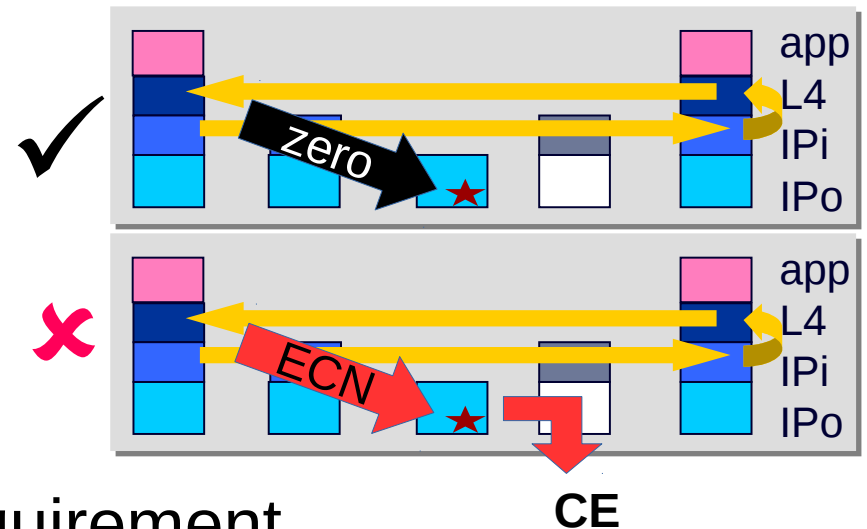
incoming inner	incoming outer			
	Not-ECT	ECT(0)	ECT(1)	CE
Not-ECT	Not-ECT	Not-ECT	Not-ECT	drop
ECT(0)	ECT(0)	ECT(0)	ECT(1)	CE
ECT(1)	ECT(1)	ECT(1)	ECT(1)	CE
CE	CE	CE	CE	CE

Outgoing header



Compliance requirement for non-RFC6040 implementations!?

- Written as an operator config requirement
 - if decap does not, or might not, propagate ECN to RFC 6040 (or equiv), if possible, the operator **MUST** configure the ingress to zero the outer ECN field



- Prerequisite implementation requirement
 - Config of ECN encap **MUST** be independent from DSCP encap
- Added text updates RFC 6040, and shim tunnel RFCs

Status and Next Steps

draft-ietf-tsvwg-rfc6040update-shim-04

- 4 revs in last IETF cycle
- Milestone: WGLC Sep 2017
- Been pushing to meet that, still feasible

- Await comments on Thu from int-area heads-up
- Teredo open issue: mtg next week to close off