

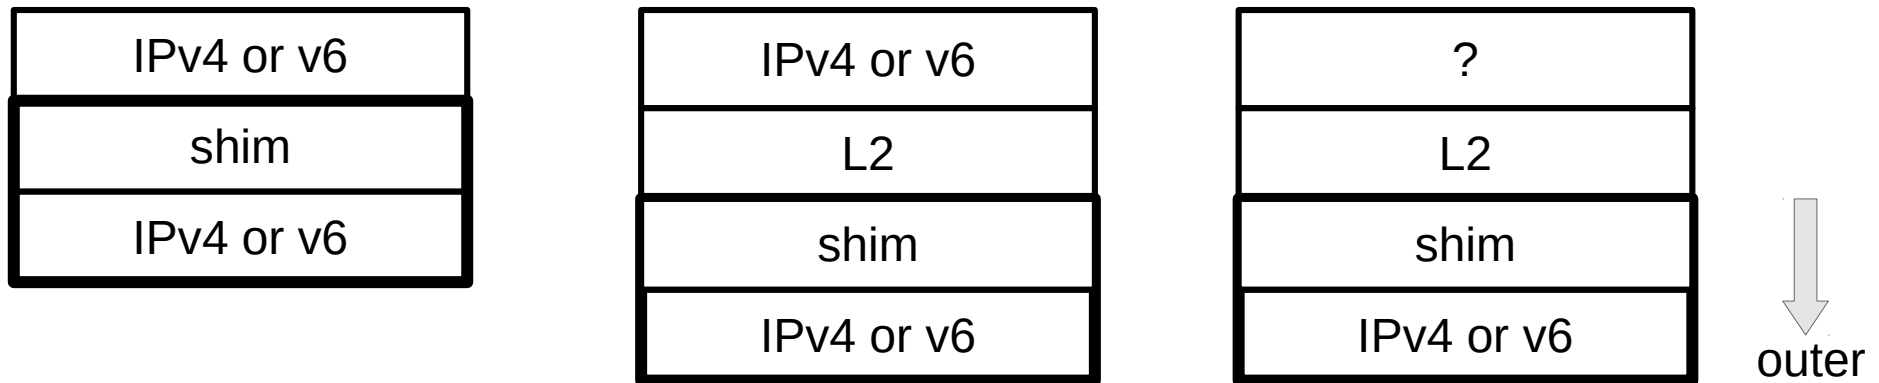
Propagating ECN across IP tunnel Headers Separated by a Shim

draft-ietf-tsvwg-rfc6040update-shim-05
(posted this morning, sorry)

Bob Briscoe
bobbriscoe.net
IETF-100, Nov 2017

Problem #1



- RFC6040 “Tunnelling of ECN”; scope was all IP-in-IP tunnels
- rfc6040update-shim 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

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				



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

Protocol	RFC	STDs or widely deployed	AOK	NOK:		NOK: non-IETF: update recom'n'd
				6040shim updates		
				Safe config	protocol	
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	✓		↑		
MIP4	5944	✓		↑	✗	
MIP6	6275	✓		↑	✗	
PMIP	5845	✓		↑	✗	
L2TPv3	3931	✓		✓	✓	
L2TPv2	2661	✓		✓	✓	
PPTP	2637	✗				
UDP	8085	✓	✓			
AMT	7450	✓		✓	✓	
TCP+IKE/IPsec	8229	✓	✓			

Why update some protocols but not others?

- In all cases, each base protocol RFC has been updated with an “operator safe configuration” clause
 - Places ref to ECN problem and solution in “Updated by” header
- Protocol Spec:
if ((a maintainer of the protocol could be found)
 && (it seemed like code might get updated)
 && (I was confident I knew the implications of the update))
 {I proposed fix to the control plane protocol, iterated, done}
else



Updates to standards track tunnel RFCs added this IETF cycle

- AMT (Automatic Multicast Tunnelling)
 - Updates to RFC 7450:
 - defined new flag on Request message for gateway to declare its ECN capability to the relay that will tunnel towards it (unidirectional)
 - Operator required to follow safe config in present spec
 - ACK: Jake Holland
- GRE (Generic Routing Encapsulation)
 - Update to RFC2784:
 - Operator required to follow safe config in present spec
 - Referred to (but not updated) 4 control protocols that are known to set up GRE tunnels
 - MIP4, MIP6, PMIP, IKEv2
 - ACK: Sri Gundavelli

Status and Next Steps

draft-ietf-tsvwg-rfc6040update-shim-05

- Rev posted this morning (sorry)
- Milestone: WGLC Sep 2017
- Been pushing to meet that, ready now - not so late
- Corridor chat this week might lead to updates to outstanding protocols (Teredo, (P)MIP)
 - Assume not

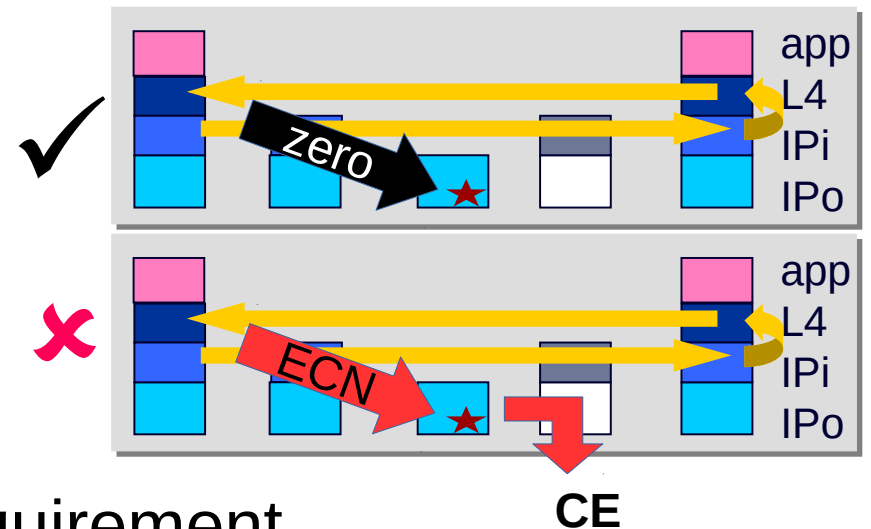
Propagating ECN across IP tunnel Headers Separated by a Shim

draft-ietf-tsvwg-rfc6040update-shim-05

Q&A

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