

Tunnelling of Explicit Congestion Notification

[draft-briscoe-tsvwg-ecn-tunnel-08.txt](#)

PCN-specific highlights

Bob Briscoe, BT
IETF-77 pcn Mar 2010




This work is partly funded by Trilogy, a research project supported by the European Community www.trilogy-project.org



status

- Tunnelling of Explicit Congestion Notification
 - **revised WG draft:** [draft-ietf-tsvwg-ecn-tunnel-08.txt](#) 03 Mar '10
 - **intended status:** standards track
 - **updates:** 3168, 4301 (if approved)
 - **RFC pub target:** Dec '09
 - **immediate intent:** in WG last call & Security Directorate review
 - **w-gs & r-gs affected:** TSVWG, PCN, ICCRG, IPsecME, Int Area?
- revised four times since last IETF, 04 - 08:
 - consensus on functional changes & alarms
 - additions for PCN support remain intact
 - tightening up of normative words
 - PCN-specific appendices marked for deletion – added summaries in main body
 - re-reviews: Gorry Fairhurst, David Black
 - new reviews: Michael Menth, Teco Boot
- minutiae are important – these are changes to IP

recap of the tunnel ingress issue



- RFC4301 IPsec supported PCN, RFC3168 ECN did not
 - multi-bottleneck excess rate marking
 - ingress hides first marking from second
- harmonise back to one branch
 - with the PCN support of IPsec

①

incoming header (also = outgoing inner)	outgoing outer		
	RFC3168 ECN limited functionality	RFC3168 ECN full functionality	RFC4301 IPsec
Not-ECT	Not-ECT	Not-ECT	Not-ECT
ECT(0)	Not-ECT	ECT(0)	ECT(0)
ECT(1)	Not-ECT	ECT(1)	ECT(1)
CE	Not-ECT	ECT(0)	CE

ecn-tunnel unchanged **compatibility mode** for legacy

'reset' CE no longer used

becomes **normal mode** for all IP in IP

changes to standards actions


draft-04 → 08

- normal mode at ingress (§4.3)
 - distinction much clearer: "MUST implement" and "SHOULD use"
 - otherwise could be lazily interpreted as "SHOULD implement"
 - if only implement compatibility mode, wouldn't add ECN/PCN support
 - closes "compliant if do nothing" loophole used in the past

Incoming Header	Outgoing Outer Header	
	Compatibility Mode	Normal Mode
Not-ECT	Not-ECT	Not-ECT
ECT(0)	Not-ECT	ECT(0)
ECT(1)	Not-ECT	ECT(1)
CE	Not-ECT	CE

recap of ingress modes

recap egress behaviour in existing RFCs



- OK for current ECN
 - 1 severity level of congestion
- any outer changes betw ECT(0/1) lost
 - reason: to restrict covert channel (but 2-bit now considered manageable)
 - effectively wastes ½ bit in IP header
- **prevents PCN using this transition**

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


Outgoing header (RFC4301 \ RFC3168)

'final' egress rules (since -05)

supports 2 severity levels of congestion marking in one DSCP
draft-ietf-pcn-3-in-1-encoding

CU but forwarded so usable in future;
still drop CE as a 'backstop';
IPsec & non-IPsec still consistent

encapsulation at tunnel ingress



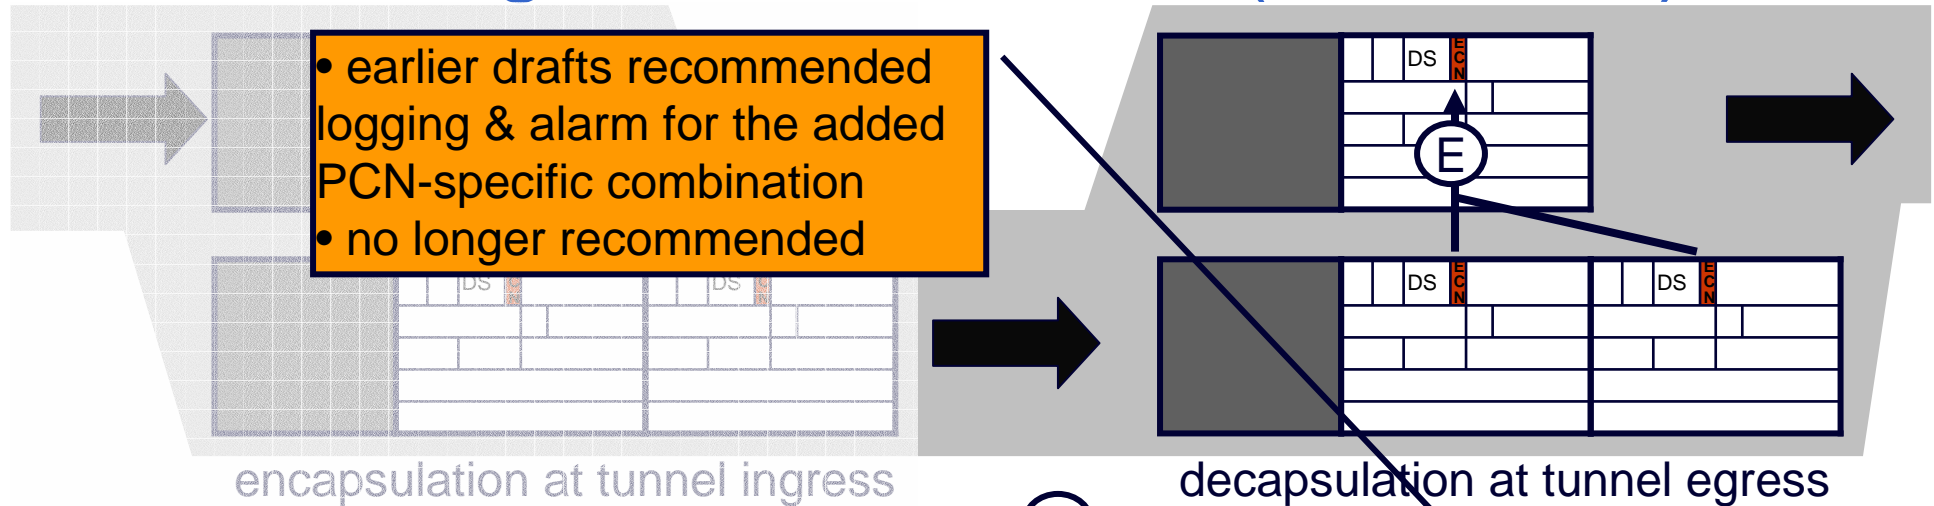
decapsulation at tunnel egress

- cater for ECT(1) meaning either more severe or same severity as ECT(0)
 - for PCN or similar schemes that signal 2 severity levels
- drop potentially unsafe unused combination
 - where high severity congestion marked in outer but inner says transport won't understand

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 (proposed update)
(bold = proposed change for all IP in IP)

'final' egress CU alarms (since -05)



- cater for ECT(1) meaning either more severe or same severity as ECT(0)
 - for PCN or similar schemes that signal 2 severity levels
- drop potentially unsafe unused combination
 - where high severity congestion marked in outer but inner says transport won't understand
- only changing currently unused combinations
 - optional alarms added to unused combinations
- only tunnels that need the new capability need to comply
 - an update, not a fork
 - no changes to combinations used by existing protocols (backward compatible)

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 (proposed update)
(bold = proposed change for all IP in IP)

- 3 types of currently unused (SHOULD log, MAY alarm)
1. (!!!) = always CU, always potentially dangerous
 2. (!) = always CU, possibly dangerous
 3. CU in this deployment (operator specific)

next steps

- In WG last call & Security Directorate review
- issues or messages of support to tsvwg list please

Tunnelling of Explicit Congestion Notification

[draft-briscoe-tsvwg-ecn-tunnel-08.txt](#)

PCN-specific highlights

