Tunnelling of Explicit Congestion Notification

draft-briscoe-tsvwg-ecn-tunnel-08.txt

Bob Briscoe, BT

IETF-77 tsvwg 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
  - tightening up of normative words
  - editorial changes – now focused & stable
  - re-reviews: Gorry Fairhurst, David Black
  - new reviews: Michael Menth, Teco Boot
- minutiae are important – these are changes to IP

recap egress behaviour in existing RFCs

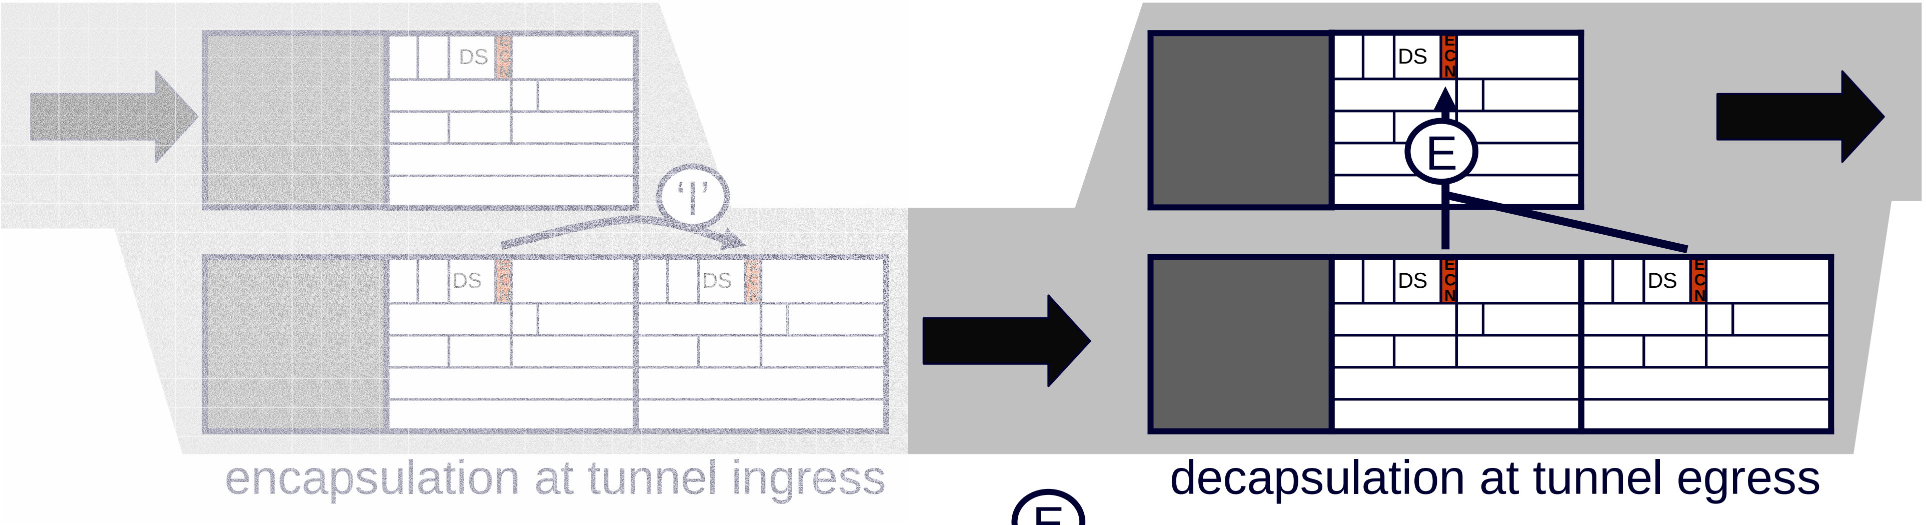
DS
ECN
'I'
E
encapsulation at tunnel ingress
decapsulation at tunnel egress

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

E

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

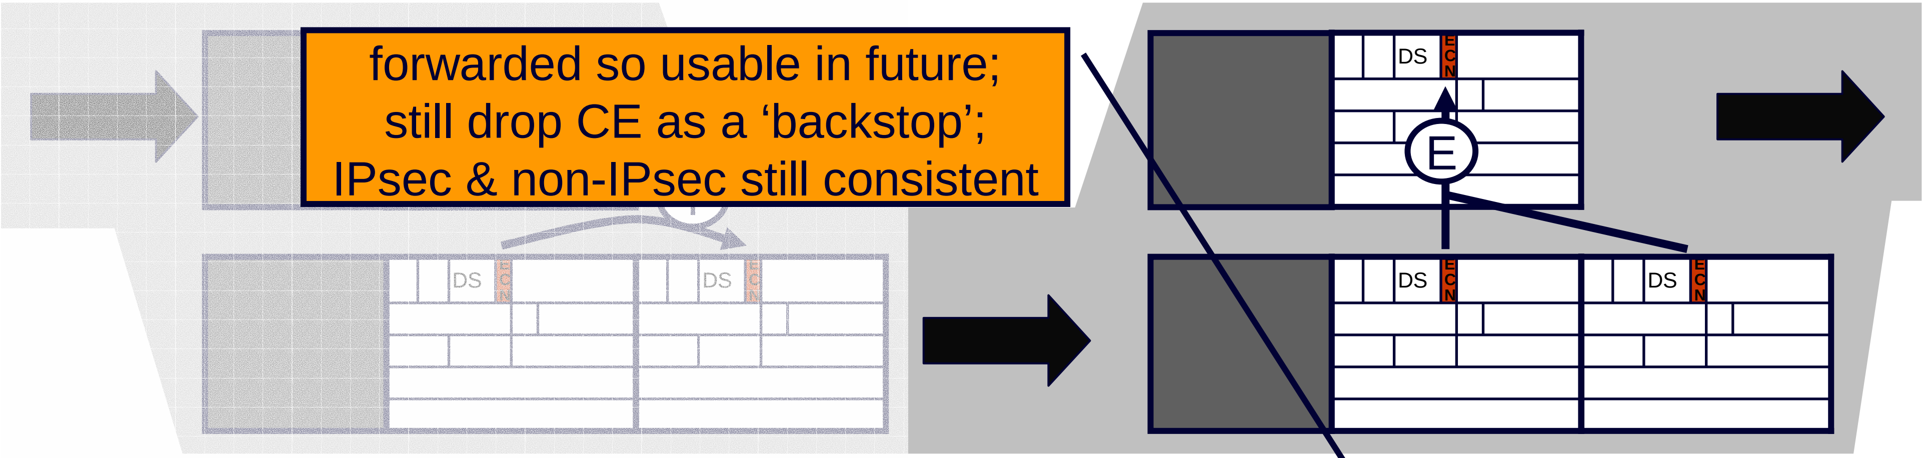

forwarded so usable in future;
still drop CE as a ‘backstop’;
IPsec & non-IPsec still consistent

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)

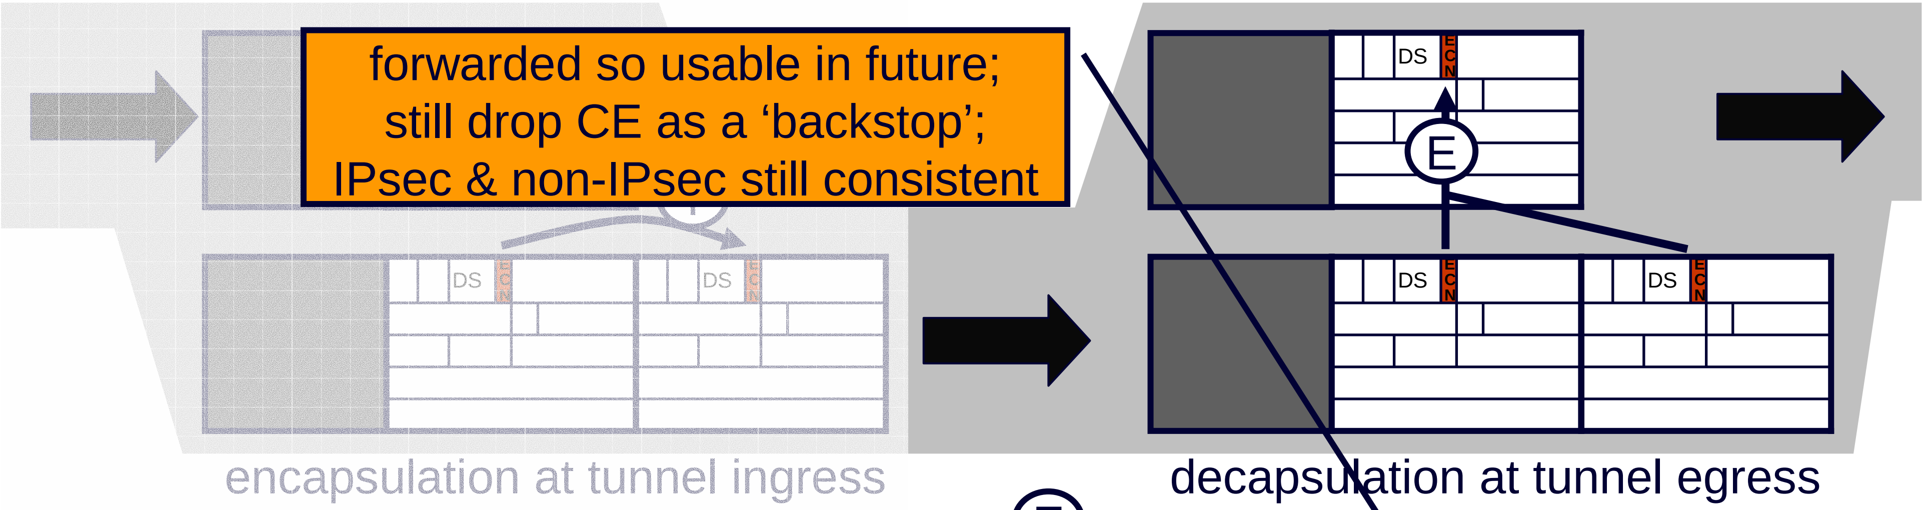

forwarded so usable in future;
still drop CE as a ‘backstop’;
IPsec & non-IPsec still consistent

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
- 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)

E

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)

changes to standards actions draft-04 → 08

- whether to design alternate ECN tunnelling (§4)
  - changed non-RFC2119 phrase 'NOT RECOMMENDED' to 'SHOULD be avoided‘
- advice on designing alternate ECN tunnelling (§7)
  - altered to reflect the functional changes (previous slide)
  - changed any upper-case keywords in the informative section to lower case.
- used upper-case in 'Alarms SHOULD be rate-limited‘ (§4.2)
- 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 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

- cut out corner-case concerning manual keying of IPsec tunnels (§5.1)
  - left as note “to be deleted by RFC Ed” during Security Directorate review

main editorial changes
draft-04 → 08

- emphasised harmonisation of fork (non-IPsec & IPsec)
  - both pre-existing branches still work as before
  - any tunnel can be deployed unilaterally without any modes or configuration
  - aim for ECN field to behave consistently whatever tunnels intervene
- altered section on updates to earlier RFCs
  - described updates to implementations, not updates to RFC text
- summarised PCN-related rationale in body
  - marked appendices giving full rationale “to be deleted by RFC Ed”
- updated acks; recent reviewers & re-reviewers
  - Teco Boot, Michael Menth, Gorry Fairhurst & David Black
- usual minor textual clarifications

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

Q&A