Google

ConEx Concepts and Abstract Mechanism

draft-mathis-conex-abstract-mech-00.txt

Matt Mathis, Google
Bob Briscoe, BT
(two duffers)
presented instead by Andrea Soppera, BT

IETF-79 ConEx Nov 2010

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

trilogy

BT

ConEx Concepts and Abstract Mechanism

- **new individual draft:** draft-mathis-conex-abstract-mech-00.txt
- **intended status:** informational
- **immediate intent:** request adoption as ConEx w-g item
- **milestone target:** Jul 2011

recall

- defer encoding to avoid obscuring underlying design
  - abstract design of algorithms & protocol
  - encoding in different protocol headers can follow (IPv6, v4)
- scope
  - loss-based (for incremental deployment), not just ECN
  - any transport, ConEx just using TCP as first concrete step

basic signals and functional units

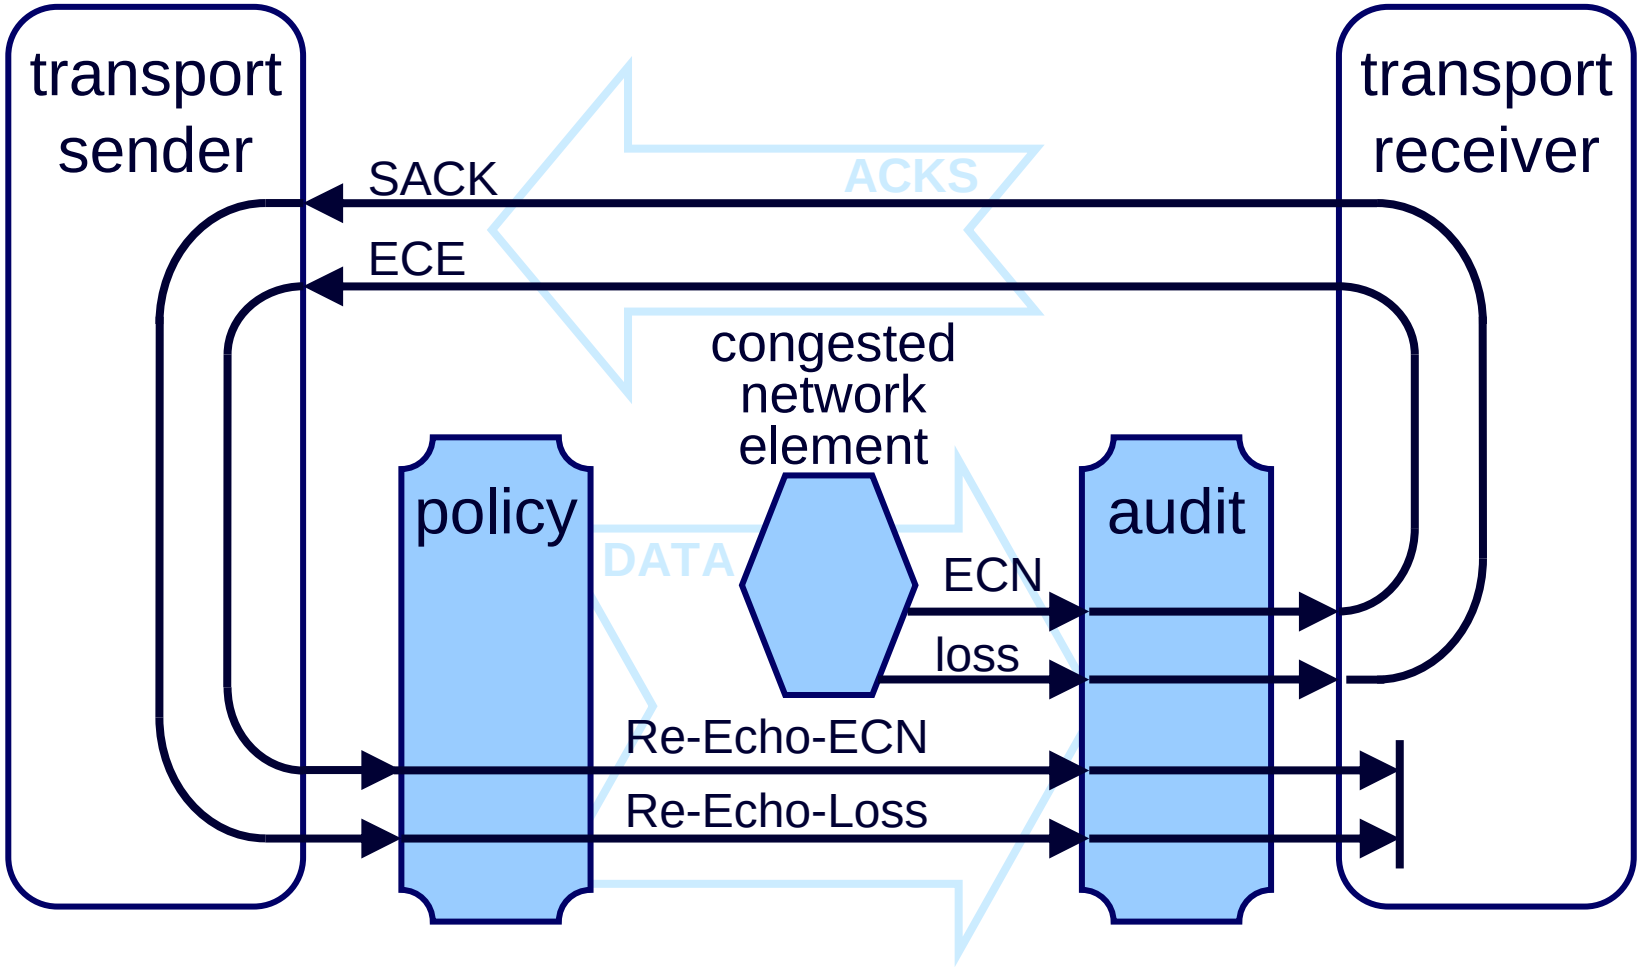

ConEx signal requirements

- visible to internetwork layer
- useful under partial deployment
  - minimal deployment: transport sender-only
- accurate (auditable)
- timely

all SHOULDs not MUSTs

in case compromises needed for encoding in headers

terminology for signalling states

max 5 states needed (white backgrounds)

as well as 3 markings, 2 types of non-marking

- Not-ConEx = ConEx not supported
- ConEx-Capable = ConEx supported
  - ConEx-Not-Marked ConEx support but not marked (yet)
  - ConEx-Marked one of:
    - Re-Echo-Loss
    - Re-Echo-ECN
    - Credit … (see later talk)

also sets of states (blue backgrounds) given names

- all names can be bashed on list

combinations

five signals do not require five flags

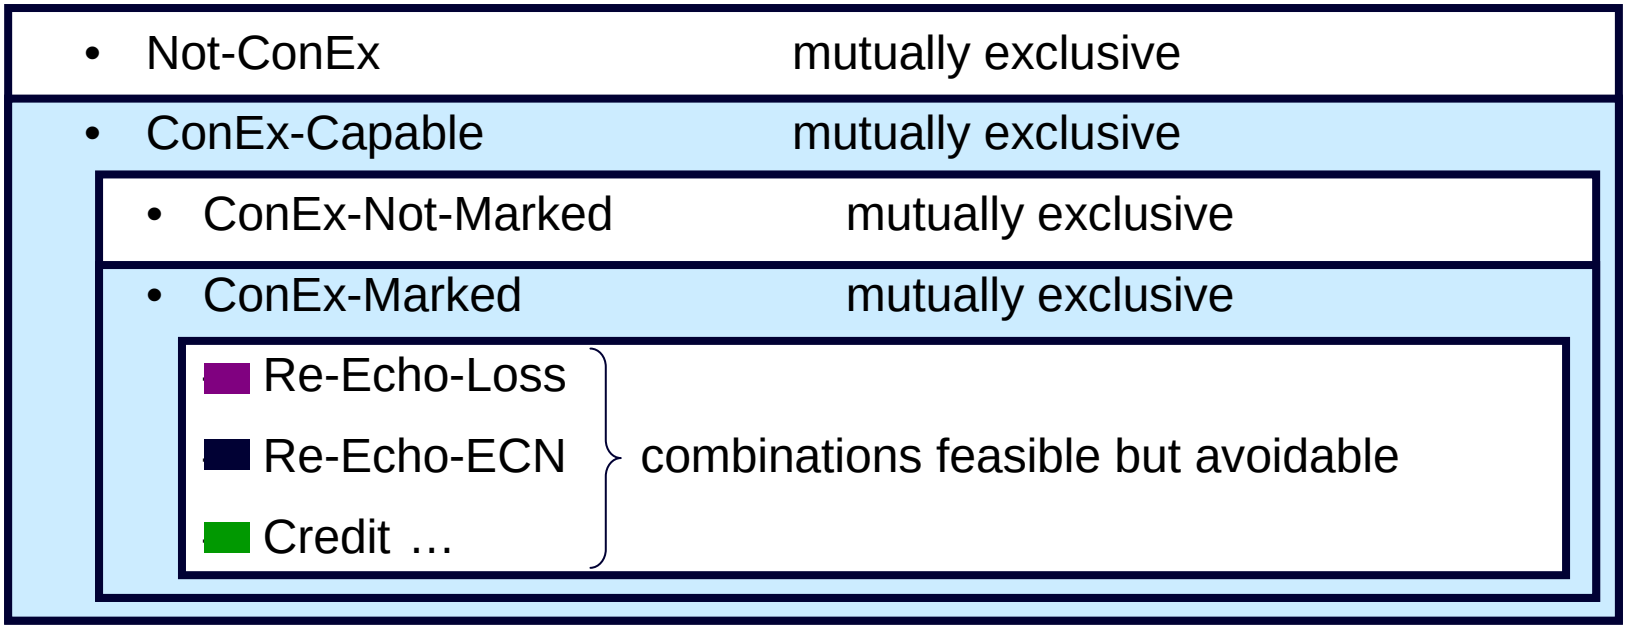

- ideally completely orthogonal to ECN
  - Re-Echo-ECN with Not-ECN-capable could be redundant
  - but may need further compromises to encode within header space

relation to re-ECN

- re-ECN: original concrete candidate ConEx proposal

  <draft-briscoe-tsvwg-re-ecn-tcp-09> up-rev'd for reference only

  implemented, security analysed
- re-ECN required ECN-capable receiver
  - *could* severely constrain deployment
- re-ECN did not need any ECN in queues
  - re-echoed loss as proposed in ConEx
  - but had no distinction between Re-Echo-ECN and Re-Echo-Loss

congestion exposure components

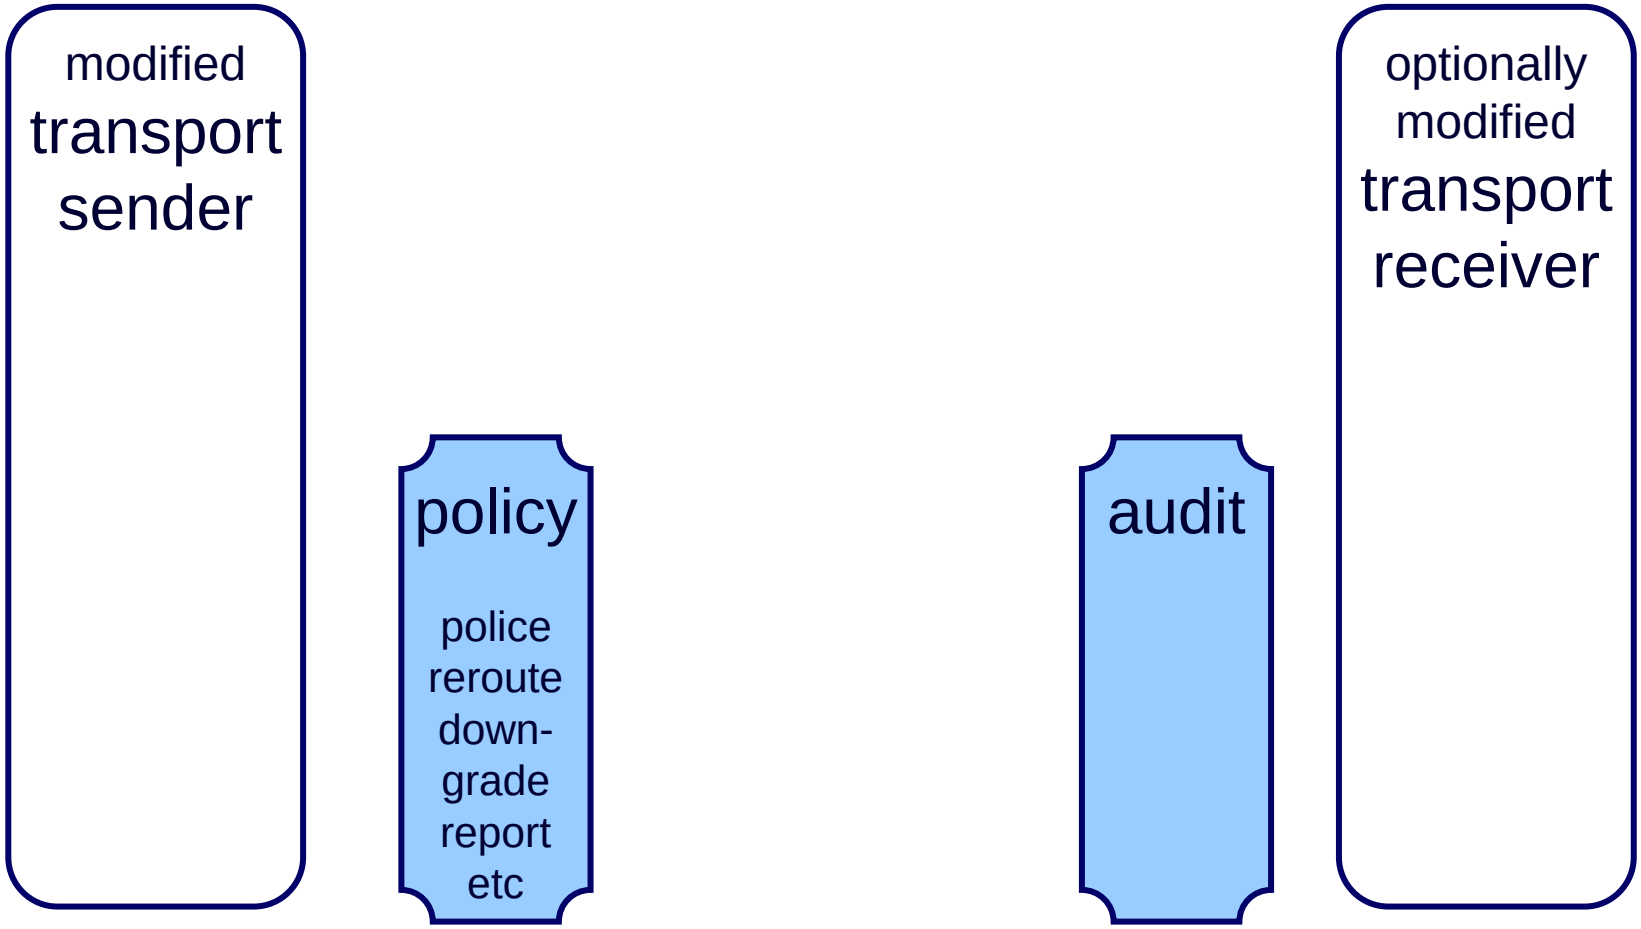

audit function

- ECN-based audit
  - counting ECN markings
  - best near receiver
- loss-based audit
  - Not a generic solution but possibly good enough in two common cases:

1. reconstruct losses by sniffing TCP seq numbers
   - Broken by IPsec, deviant TCPs
2. single primary access bottleneck
   - Bottleneck device can also perform audit

status & plans

- 5 reviews on list so far – 1 more detailed [Bagnulo]
  - all agree Credit needs to be explained (see later presentation)
  - other places where too much reader knowledge assumed
  - fairly easy to fix
- plans
  - consensus on terminology (list)
  - text to explain Credit & reach consensus if disagreement
  - add normative design criteria for audit function
  - otherwise, looking in fairly good shape
- adopt as WG draft?

ConEx Concepts and Abstract Mechanism

draft-mathis-conex-abstract-mech-00.txt

Q&A