DoS-resistant Internet - progress

Bob Briscoe

Jun 2005

BT activity

- Research
  - 2020 Communications Architecture project
    - DoS-resistant Internet Architecture task
  - Network Security project
    - BGP security
    - control plane separation
    - intrusion detection systems
- Engineering
  - Network design
  - Second line support for operations
- Operations
  - Deployment and operation of attack mitigation technology

DoS resistant Internet architecture
BT 0506 deliverables

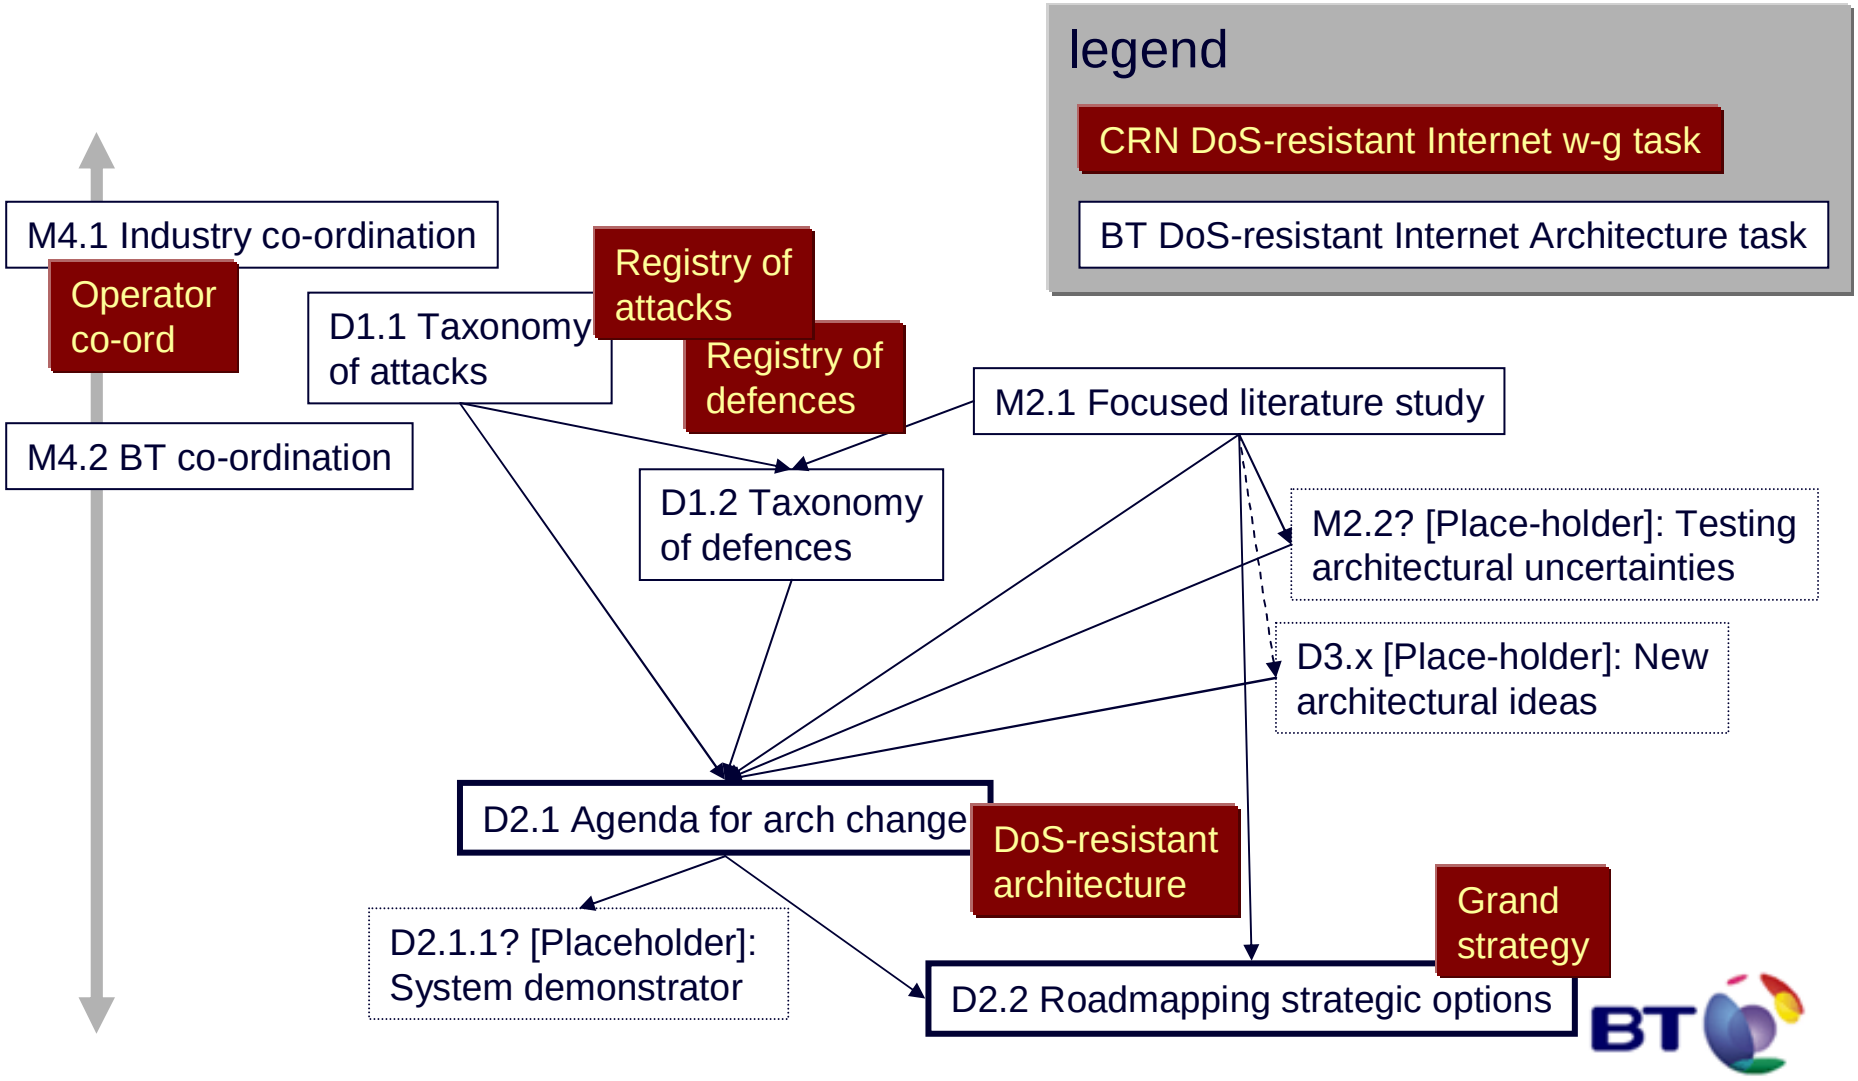

DoS-resistant Internet Architecture

- approach
  - cherry pick the ideas of others
  - sprinkle in a few ideas of our own
  - stress-test
  - propose a target architecture of complementary solutions
  - describe incremental deployment

BT

architectural component ideas
candidate list

- Symmetric paths, address separation, RPF checks, state set-up bit, nonce exchange, middlewalls
  - M Handley and A Greenhalgh “Steps towards a DoS-resistant Internet architecture” FDNA (2004)
- Secure Internet Indirection Infrastructure
  - D Adkins et al “Towards a More Functional and Secure Network Infrastructure” UC Berkeley TR-CSD-03-1242 (2003)
- Re-feedback
  - B Briscoe et al “Policing Congestion Response in an Internetwork using Re-feedback” SIGCOMM (2005)
- Receiver-driven Capabilities
  - X Yang et al, “DoS-limiting Internet architecture” SIGCOMM (2005)

- tactical approaches
  - ingress filtering, filter pushback…

symmetric paths

- powerful approach
- loss of Internet flexibility acknowledged
- extended to preserve data in flight during reroutes
- stress-testing with authors
  - big question: would it significantly reduce worm attacks?

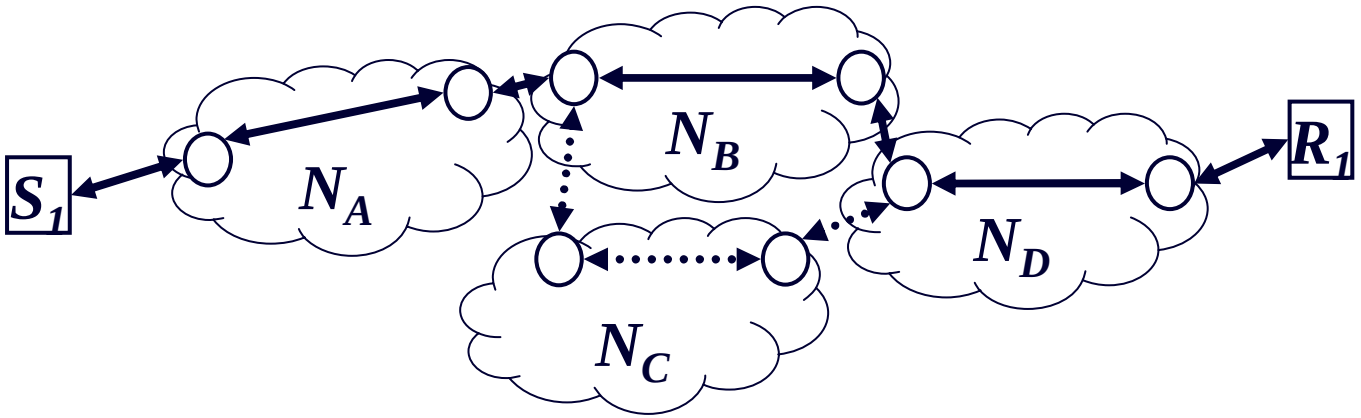

Secure Internet Indirection Infrastructure

Secure i^3

- rough analogy: receiver-driven multicast
  - receiver creates channel (trigger) in infrastructure
  - senders send to channel
- unlike IP multicast, overlay infrastructure (Chord)
  - highly redundant
- essentially, allow more, less efficient routes
  - choice of routes under receiver control
  - if some routes used for attack, drop them
  - efficient route could be norm, then less efficient routes when under attack
- inherent weakness for servers: must advertise triggers
  - so attackers on dropped triggers just re-start
  - authors offer some mitigation

BT

re-feedback incentive architecture

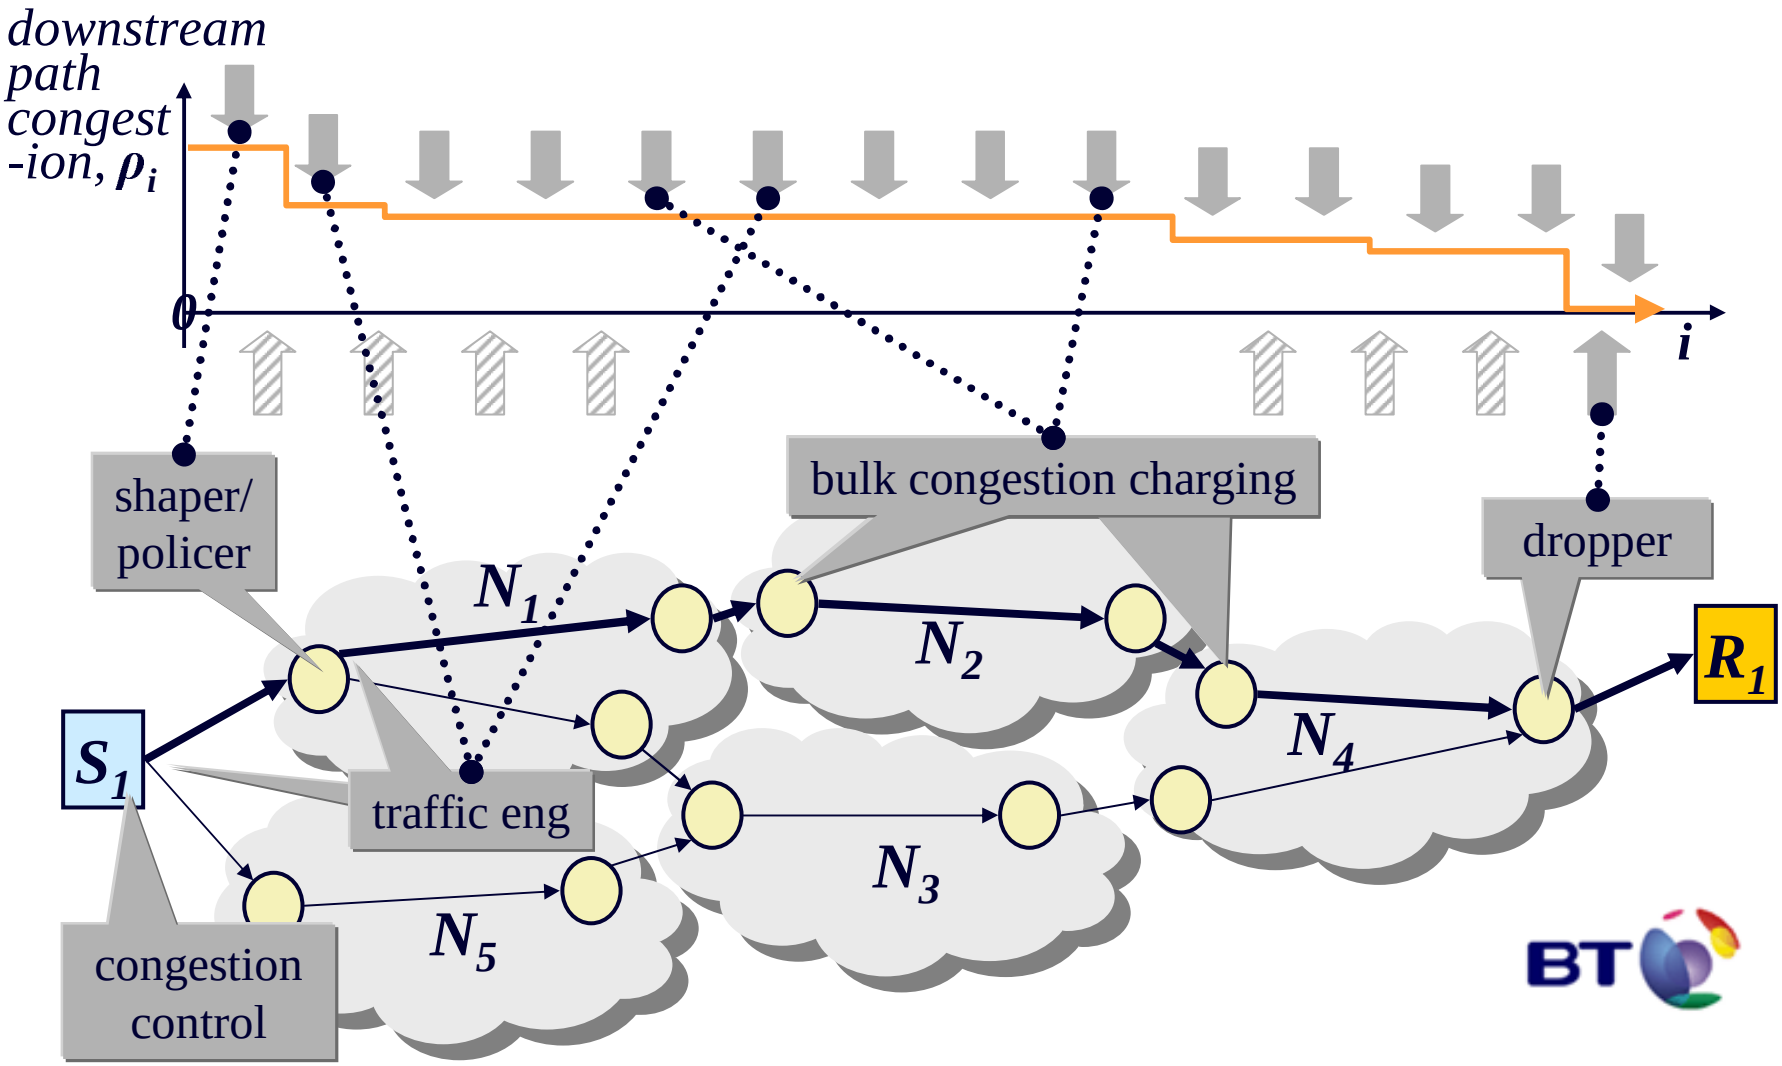

receiver-driven capabilities

- yet to fully analyse (only just published)
- sent traffic picks up time-bounded tags
  - tags from each network (router)
  - and byte permission from destination
  - collectively termed a capability
- routers store tags
- subsequent traffic authorised using capability
- detail devils
    - bootstrapping
    - bounded router state
    - incremental deployment

BT

Grand Strategy: some questions

- if upstream network doesn't filter/throttle
  - once attackers identified, what do we do?
    - continue to add more and more filters at borders?
    - disconnect their network?
    - throttle their network?
    - sue them (under what law – tort, criminal)?
- can the network help identify persistent attackers?
    - unenforceable due to numerous weak legal systems?
    - pair-wise network agreements, or source identification?
- inter-domain charging
  - congestion-based
    - would it slowly mitigate persistent attacks?
  - filter-based
    - would it encourage push-back?
- incremental deployment
    - new, clean Internet?
    - gradually clean up the one we've got

BT

in summary

- multiple answers, defence in depth
  - pair-wise network agreements AND source identification
- complementary approaches
  - identify attackers (networks) by address
  - routers filter traffic from identified attacks/attackers
  - inter-domain charge to congestion-causing networks
  - police congestion-causing traffic

BT

more info

- Bob Briscoe
- bob.briscoe@bt.com
- http://www.cs.ucl.ac.uk/staff/B.Briscoe/