


# WP3 Task 3.2 - Internal Commercial Evaluation

Research issues


19/6/2008

## TCP & sequential edge-2-edge traffic



- end-2-end traffic can sometimes be changed into a sequence of edge-2-edge traffic
  - A middle-box, or an Ingress Tunnel Router (ITR) in LISP, could be an overlay node
- How [TCP congestion control algorithm](#) should be altered so that source adapts her rate according to what happens after Overlay node A?
- How to incentivise an ISP to [pass ACKs](#) truthfully to source?
- What if ReECN is turned on and overlay nodes are 'charged' for edge-to-edge traffic they send on behalf of source?

## TCPm & sequential edge-2-edge traffic



- Maybe an ISP itself could run TCPm in an edge-2-edge way
  - How [TCPm congestion control algorithm](#) should be designed?
    - i.e. slow start as normal, could it slow start less aggressively, ...?
  - Should we incentivise inter-domain multi-path?
    - When would this be attractive? i.e. is there an upper threshold on the number of parallel paths?

## Multi-path routing & Congestion control

- Assume [HIP](#), TCPm and ReECN are in place
- How an un-trusted rendezvous server (RVS) should be incentivised to [reply with the actual \(complete\) list of Loc-ID pairs](#)?
  - RVS may have a vested interest in promoting a subset of Locators
  - A [destination](#) host has no prior information about RVS's routing policy during selection
    - In an extreme case, she may be unable to identify an opportunistic behaviour even after subscription
- How a source host should set up paths to destination given that she has not accurate information about current path 'cost' and each choice may have a setup 'cost' (i.e. slow start rate)
  - How many/which of the available Locators to use to minimize total 'cost'? When to swap to a new Locator?
    - What if the paths to different Locators have many links in common?
    - Short flows = free riders?

## Multi-path routing & Congestion control

- Assume LISP, TCPm and ReECN are in place
- How an un-trusted ITR should be incentivised to follow host's preferences? (assuming this functionality is supported...)
  - Source currently is unable to observe ITR's choices

## Impact to main stakeholders

- Users
  - Better throughput for applications?
    - Avoiding the switching cost (got to realise first path is bad, got TCP slow start on 2nd path, etc)
  - Cost uncertainty
    - If ReECN is turned on and users are somehow 'charged', then what type of network features would protect them from malware that act as a traditional (dialup) dialer?
- Network providers
  - Is a network more susceptible to attacks due to multipath routing?
    - i.e. 'malware' would not react on congestion (inelastic traffic)

## Impact to main stakeholders (cont.)

- Network providers (cont.)
  - Can multi-path congestion control do better traffic engineering than network providers can achieve today?
  - What sorts of traffic engineering games the providers would play?
    - Are providers ever motivated to try to carry more traffic? Less traffic?
      - What if ReECN is turned on?
      - Assuming TCPm, can ISP slow down the forwarding rate?
    - Does it contribute to smaller routing tables due to less need for prefix de-aggregation?

## Impact to main stakeholders (cont.)

- Network providers (cont.)
  - Multihoming should increase competition between ISPs, and multipath in general should increase competition between transit network providers.
    - How does it affect contracts?
    - How does it affect path diversity?
      - Bigger problem for network-based Loc-ID split (like LISP)?
  - Do network providers have the incentive to upgrade their network?
    - A user's traffic rate on a path signals her preferences
    - But an ISP receives information only when packets flow through her network, since the user may be silent

end