

Host to Network Signaling Use Cases for Collaborative Traffic Differentiation

draft-bwbr-tsvwg-signaling-use-cases

IETF119 — March 2024 — Brisbane

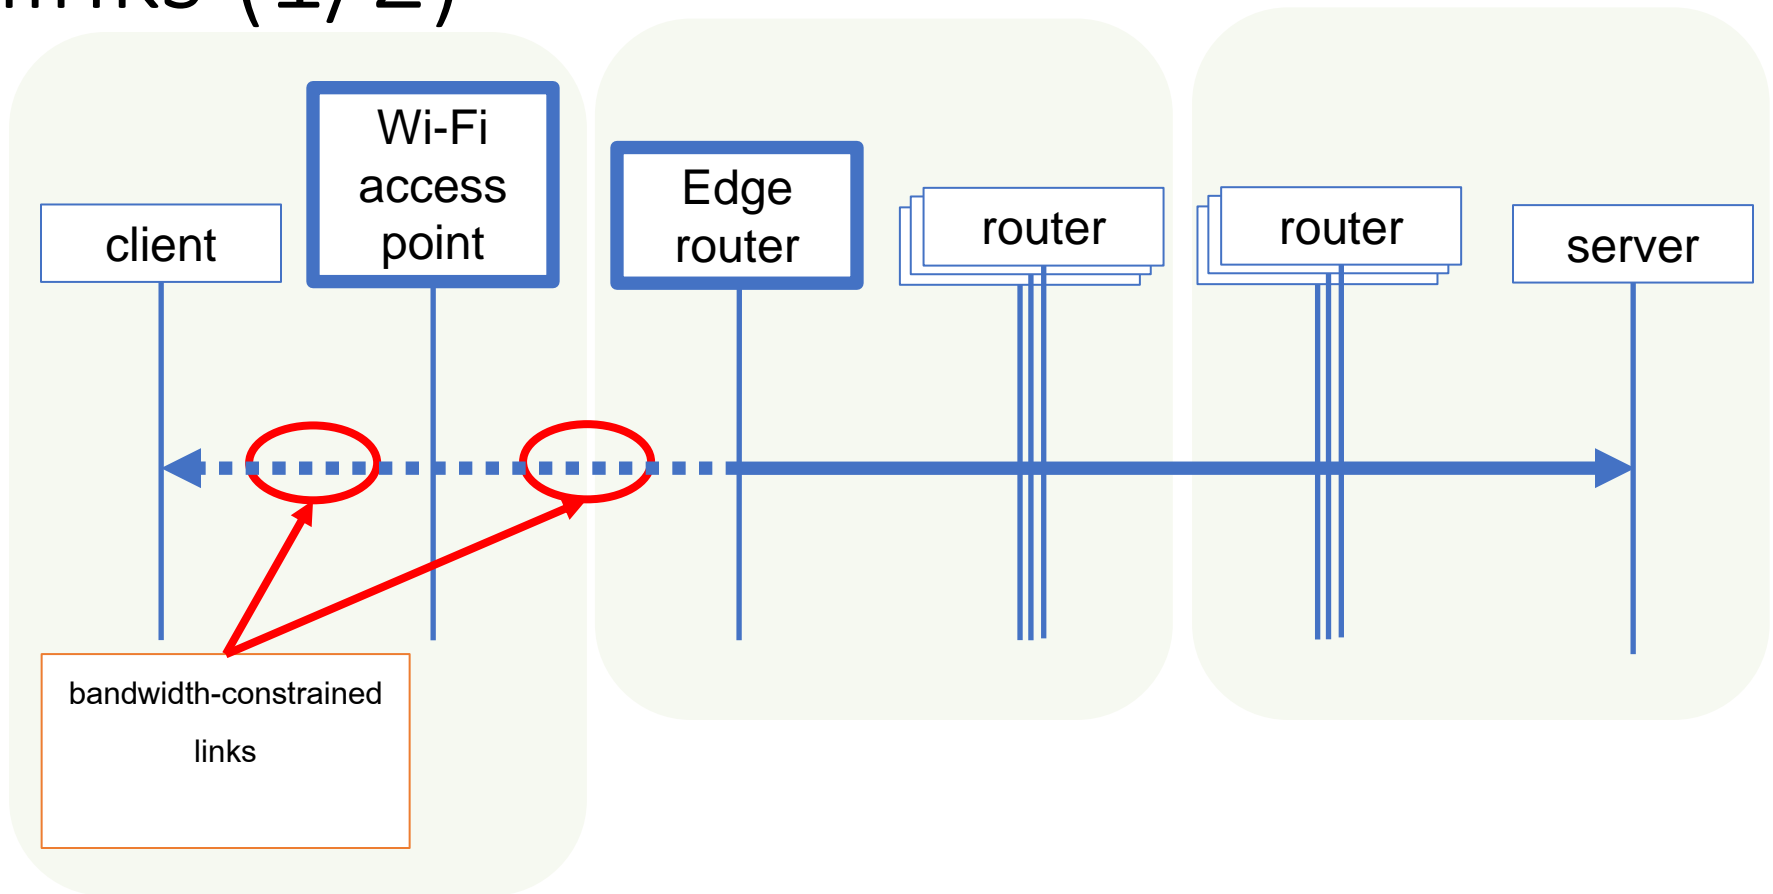
Sridharan Rajagopalan, Citrix

Dan Wing, Citrix

Mohamed Boucadair, Orange

Tiru Reddy, Nokia

Problem: bandwidth-constrained links (1/2)



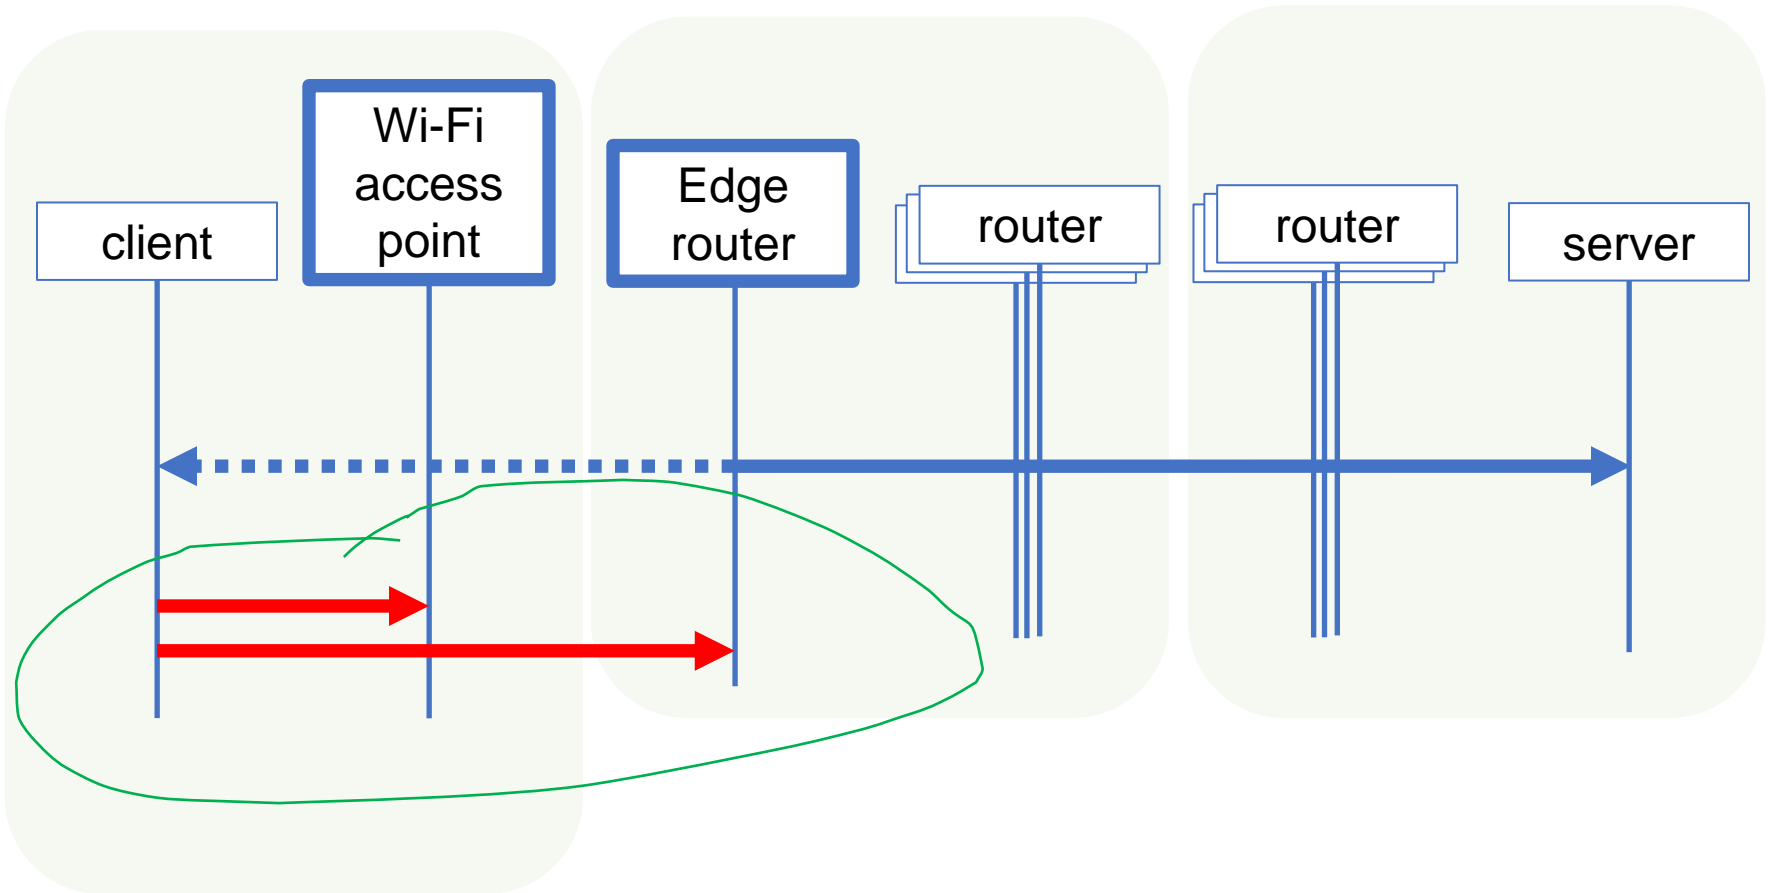
Problem: client wants to influence **incoming** packet treatment (2/2)

- Networks **cannot honor every** flow's signals
 - “Everyone Is a Fire Engine”
- Networks **cannot partner with every** sender
 - Too many senders across the Internet
- **Deployability** Tradeoffs
 - Cost vs. Benefit
 - Impact on operations vs incentive to deploy
 - Enhanced experience vs. impacts on nominal mode

High Level Requirement

- Receiver (client) *tells network which flows* have metadata
 - Client authorizes honoring incoming markings
- De-prioritize some packets of a single flow
 - *Not between flows*
 - During *reactive policy event*
- Other requirements are less important
 - Version and capability negotiation with network
- While taking into account *operational considerations to increase deployability*

Client-to-Network Signal



Where Traffic Differentiation Is Implemented?

- The network: e.g., packet discard preference
- The host: e.g., adaptive transmission or session migration
- Through cooperation of both the host and the network

What's Not Required?

- Reveal the application identity
- Reveal the server identity
- Disclose the application reason to signal metadata
- Inspect client-to-server encrypted payload by network elements

On Operational Considerations

- Some use cases suggest that networks return a maximum bitrate for a flow, however:
 - Networks usually don't assign quota per flow, application. So, *which information to return?*
 - Access control/rate limits are bound to a subscriber; not specific device. So, *how a per-host information can be computed?*
 - *Returning a global limit in a flow may be misleading*, if no scope is provided
 - Subscribers having distinct service offerings are serviced via same access node. *System-wide limits are not thus always possible*
 - Cascaded bottlenecks may be observed, *supplying a local information may not reflect the actual properties of a path*
- Clarifying intended use cases is thus key

How To Assess the Benefits?

- Have first a *common understanding* of the use cases, their requirements, and claimed expected benefit
 - Will ease running experiments and provide evidence
- However, the same signal may be shared using a variety of means:
 - How to identify potential failure roots: signal, channel, or both?
 - How to ease interconnecting domains that use distinct signaling (e.g., RAN vs. Core Network or LAN vs. Access Network)?
 - Need to *separate the channel used vs. signal itself*

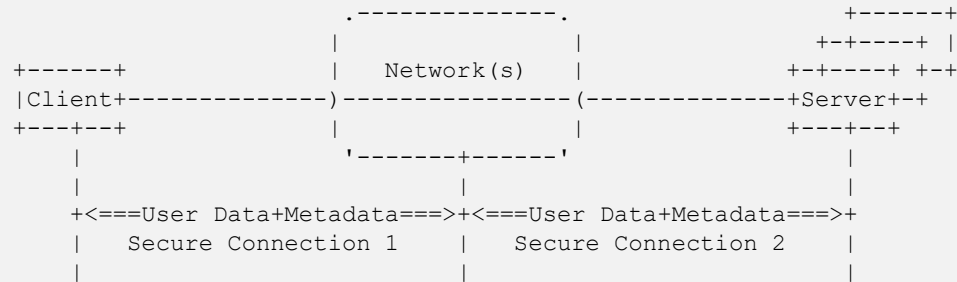
Discussion

- Do we need a requirements document for signaling?
 - Network to host
 - Host to Network

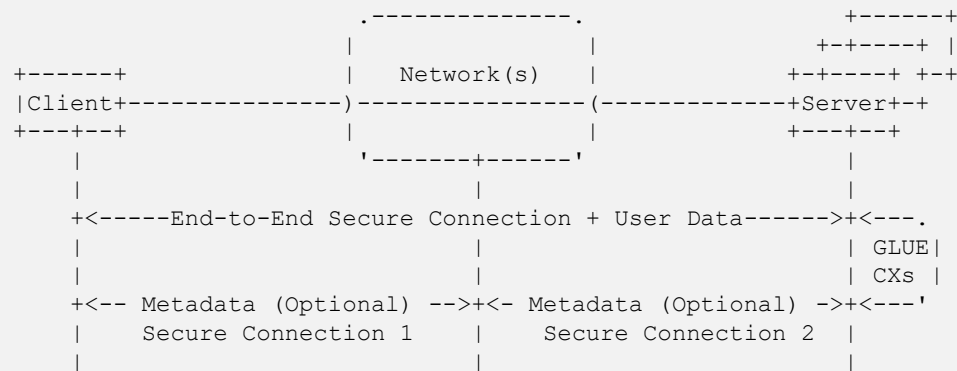
Appendix

Sample Signaling Schemes

(1) Proxied Connection



(2) Out-of-band Metadata Sharing



(3) Client-centric Metadata Sharing

