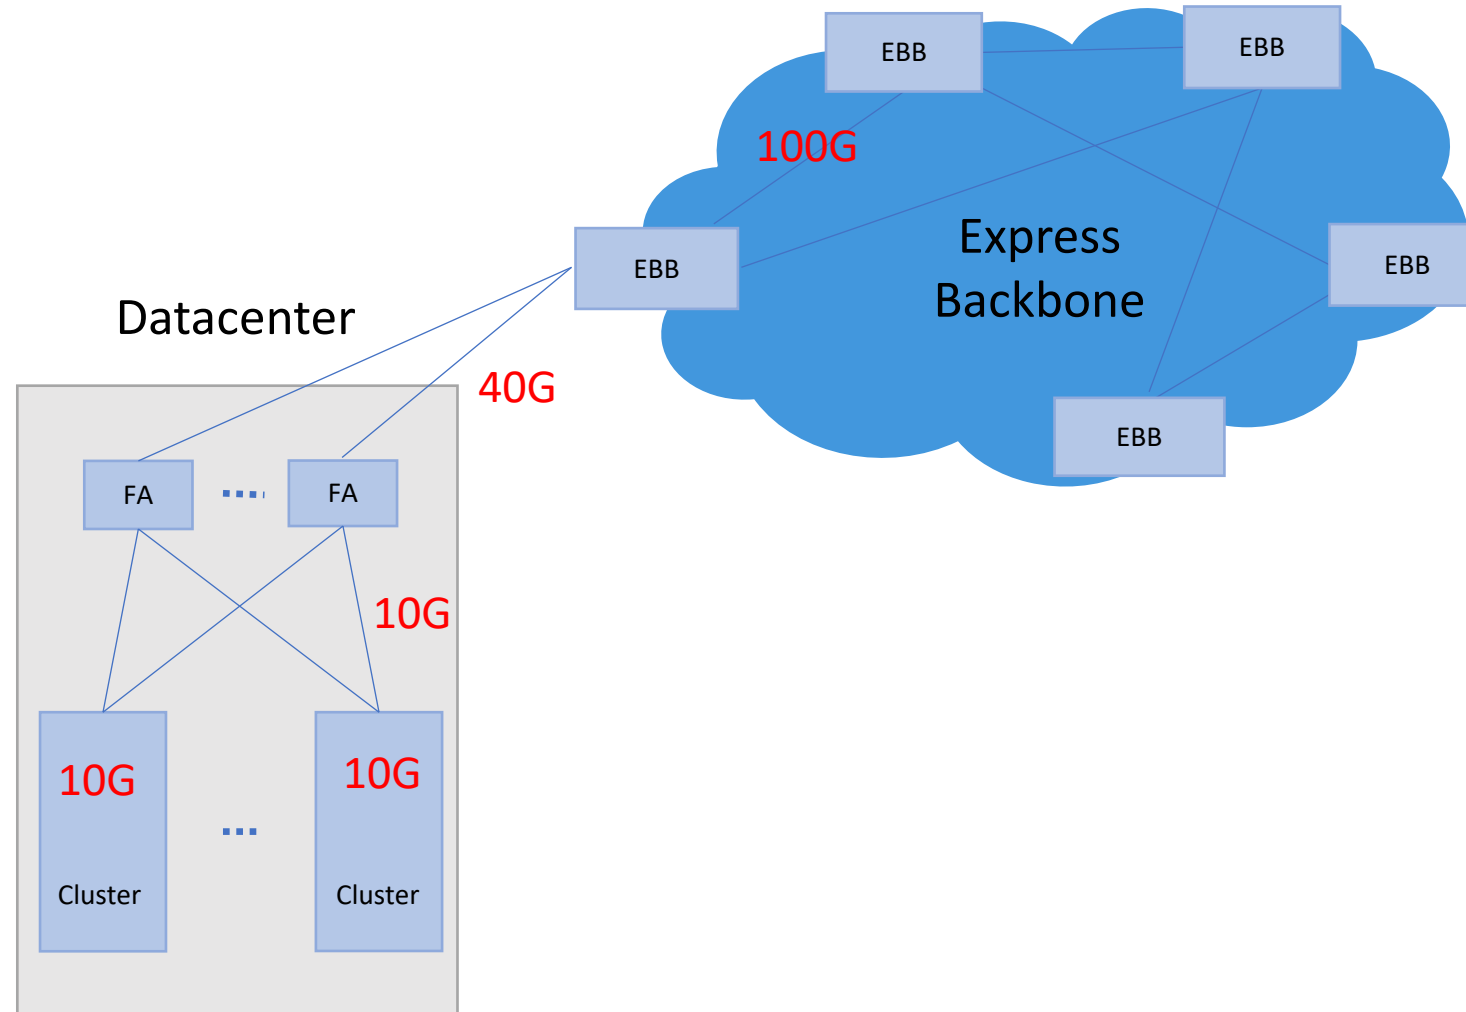


# Buffer Sizing Experiments @Facebook

2015-2016

# Routers we experimented with

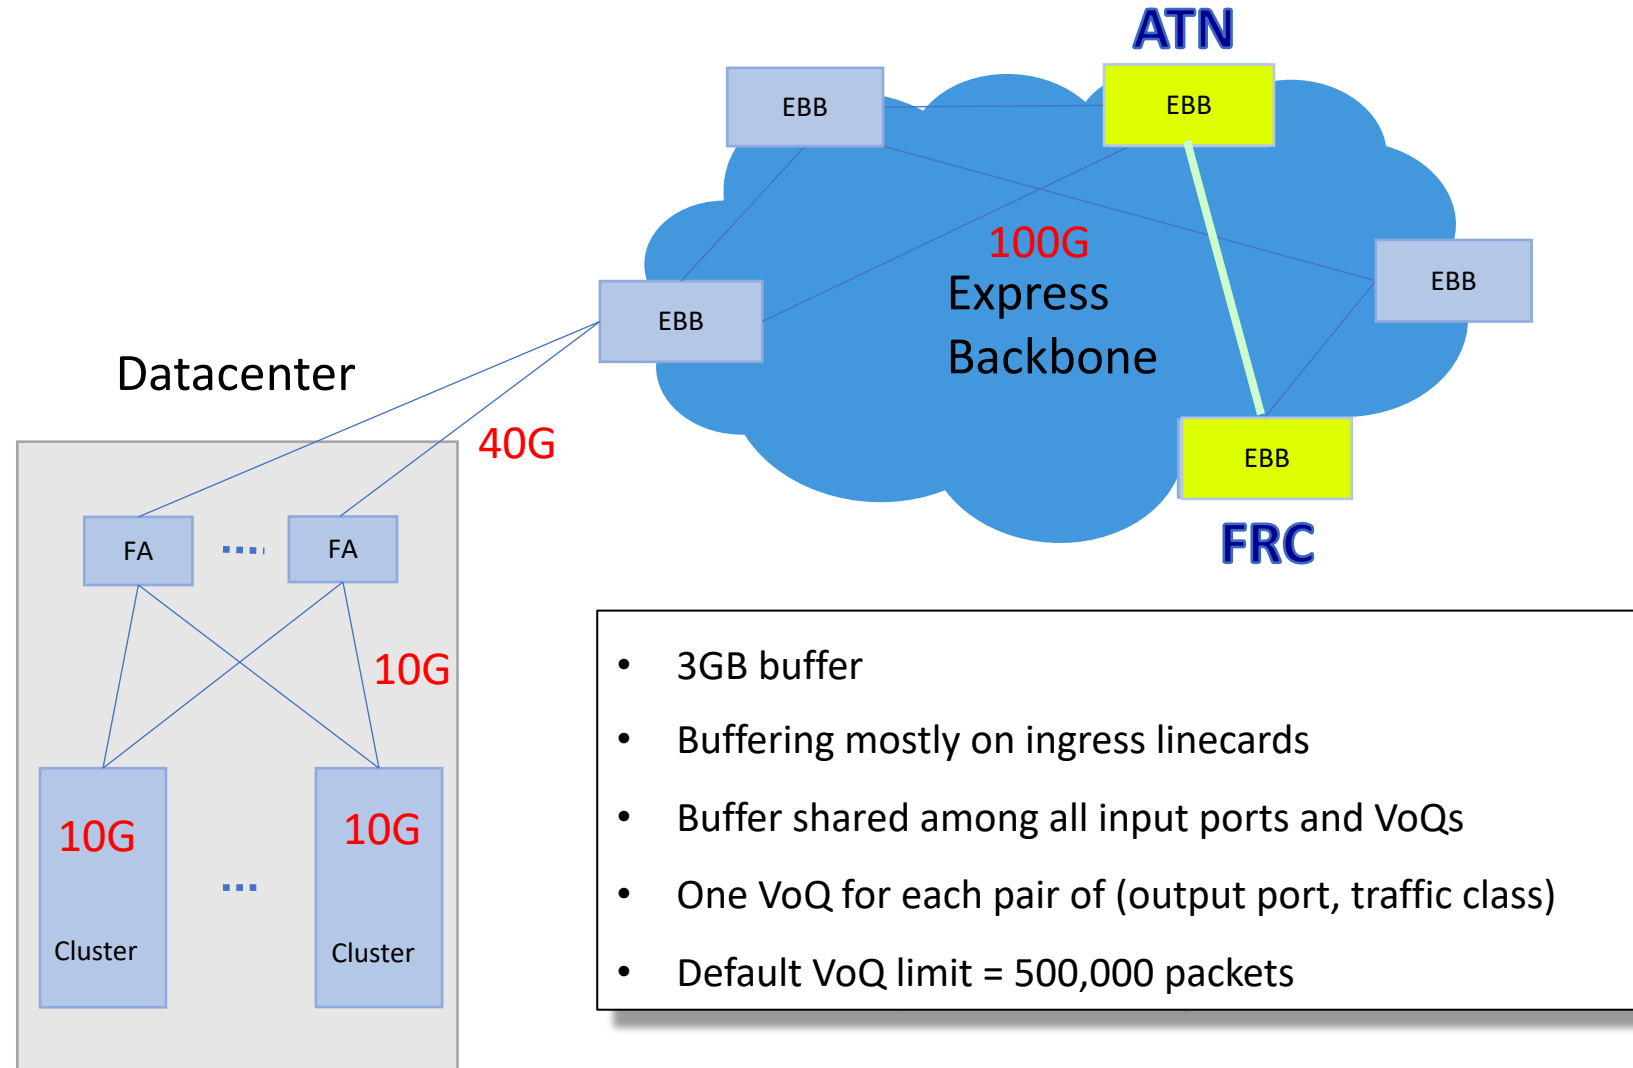


<https://code.fb.com/data-center-engineering>

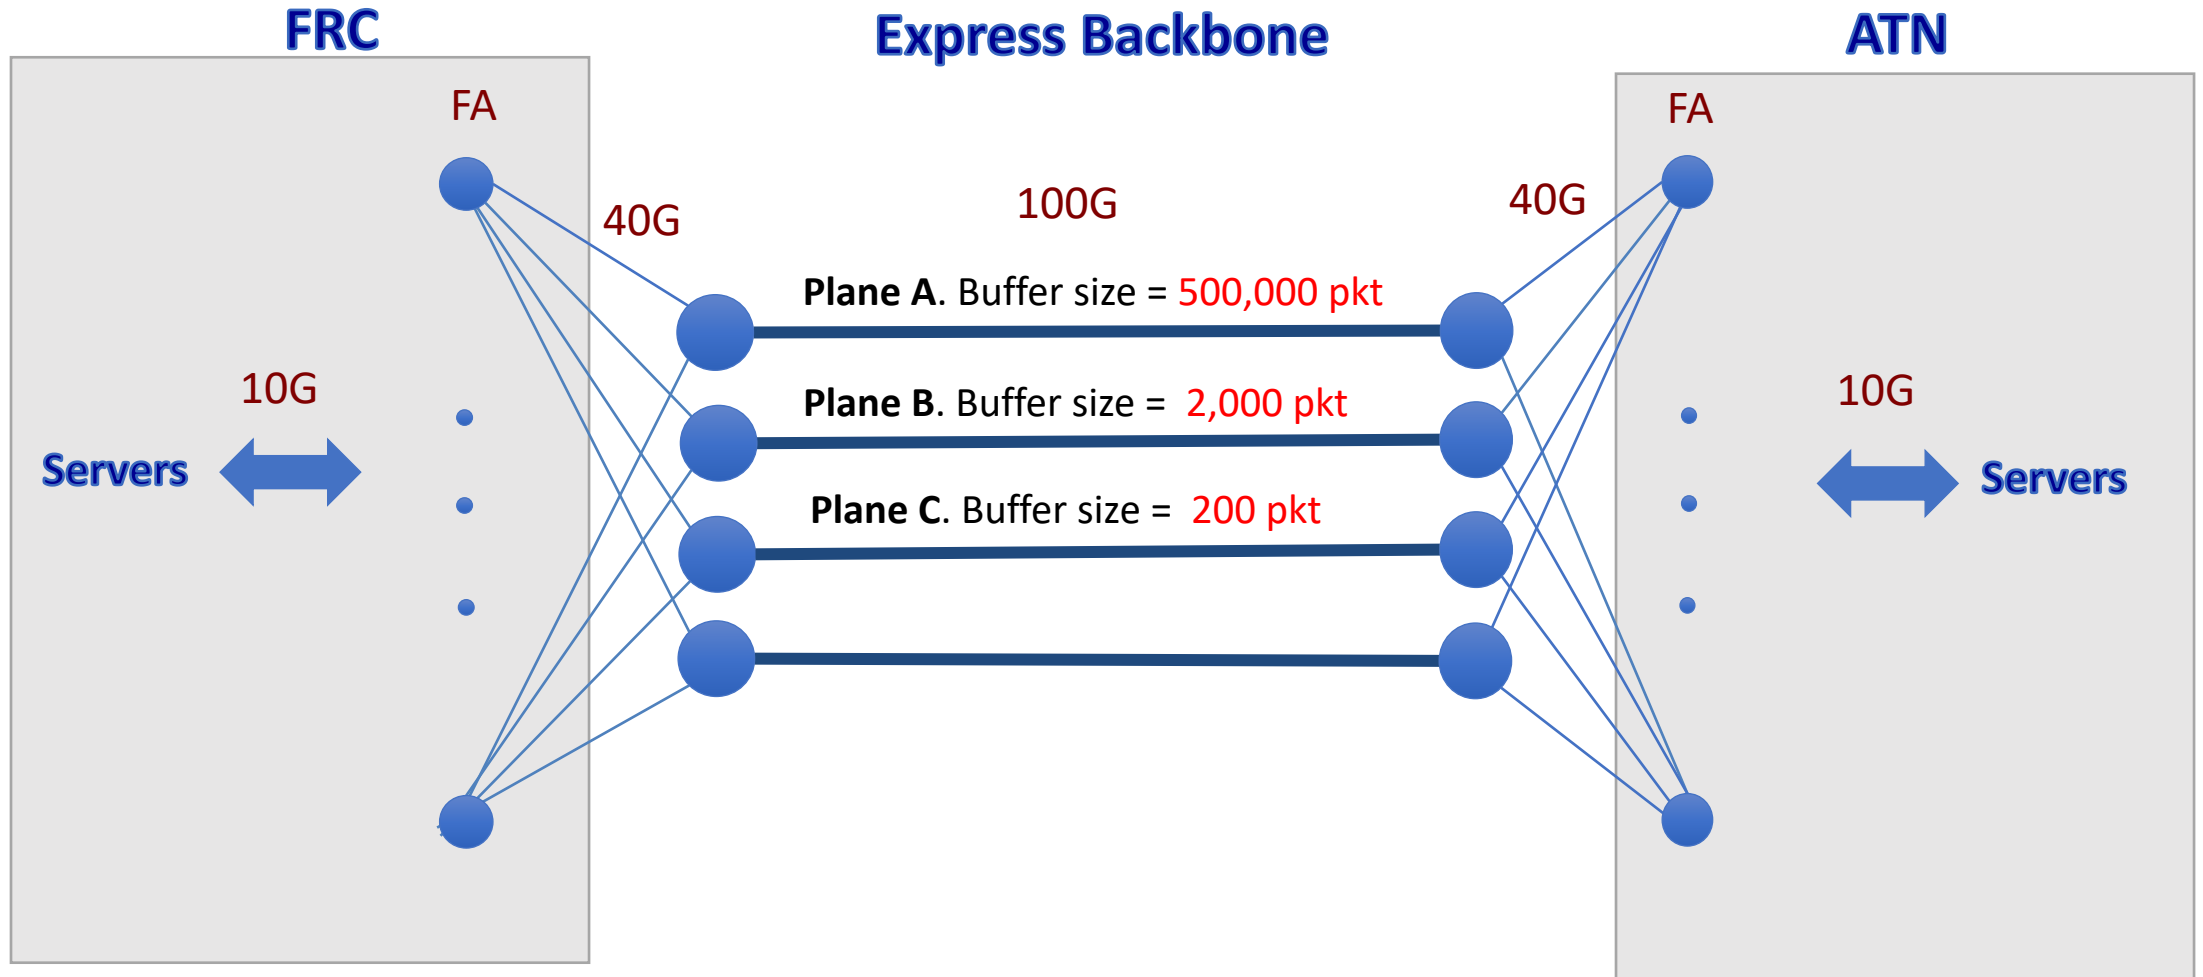
# Experiment setting

- TCP: Reno
  - No ECN. No pacing. No limit on congestion window size.
- Traffic shaping using Linux Traffic Control (TC) on some hosts
- Data Collection:
  - Coarse-grained: mostly aggregated over a few minutes
    - SNMP and our collection pipelines
    - direct collection on boxes

# Experiments in the Backbone



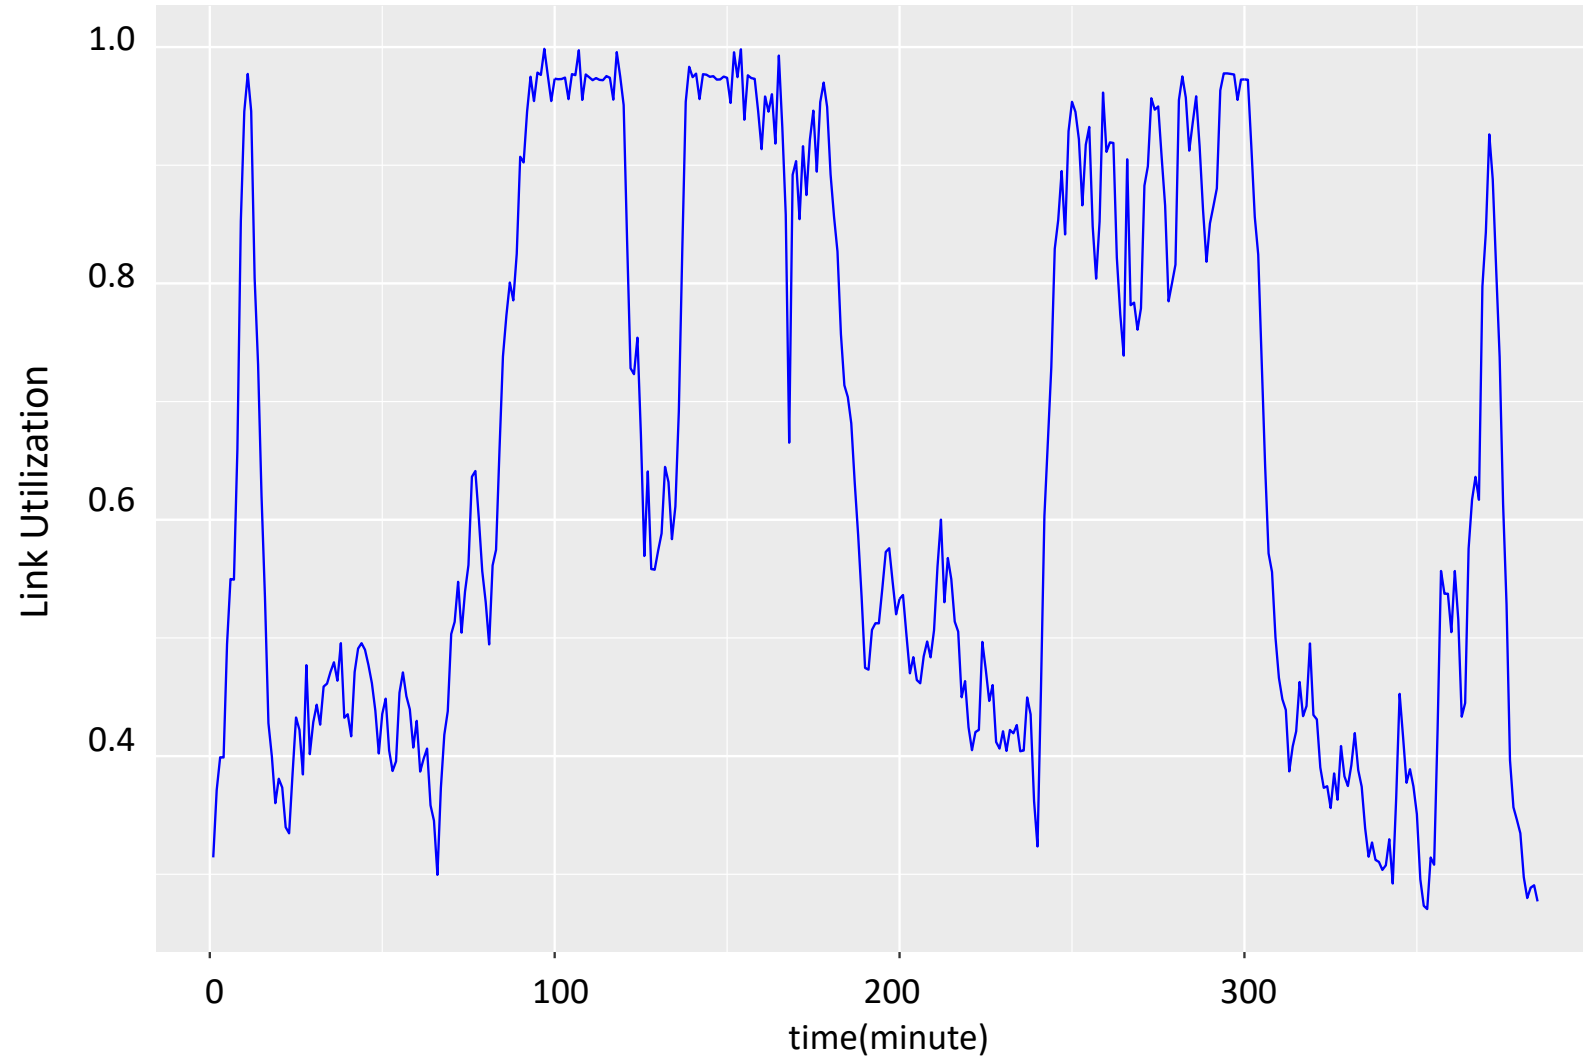
# Experiments in the Backbone



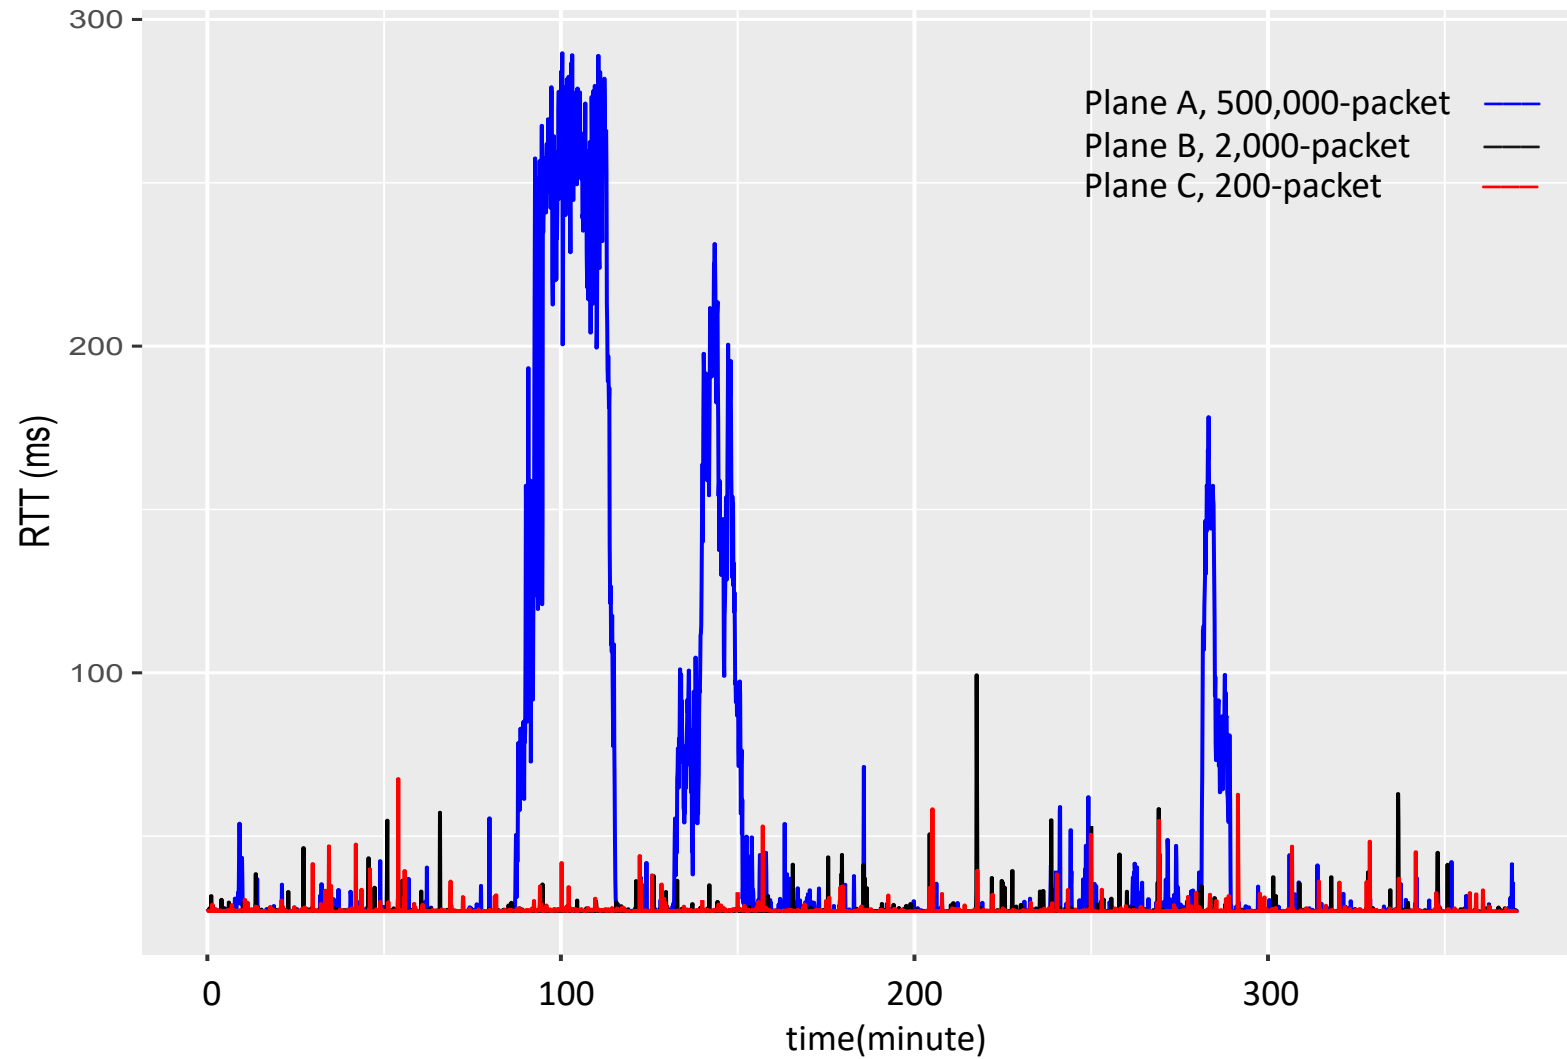
# Experiments in the Backbone

- How does buffer size affect
  1. drop rate
  2. link utilization
  3. individual flows

# Link utilization

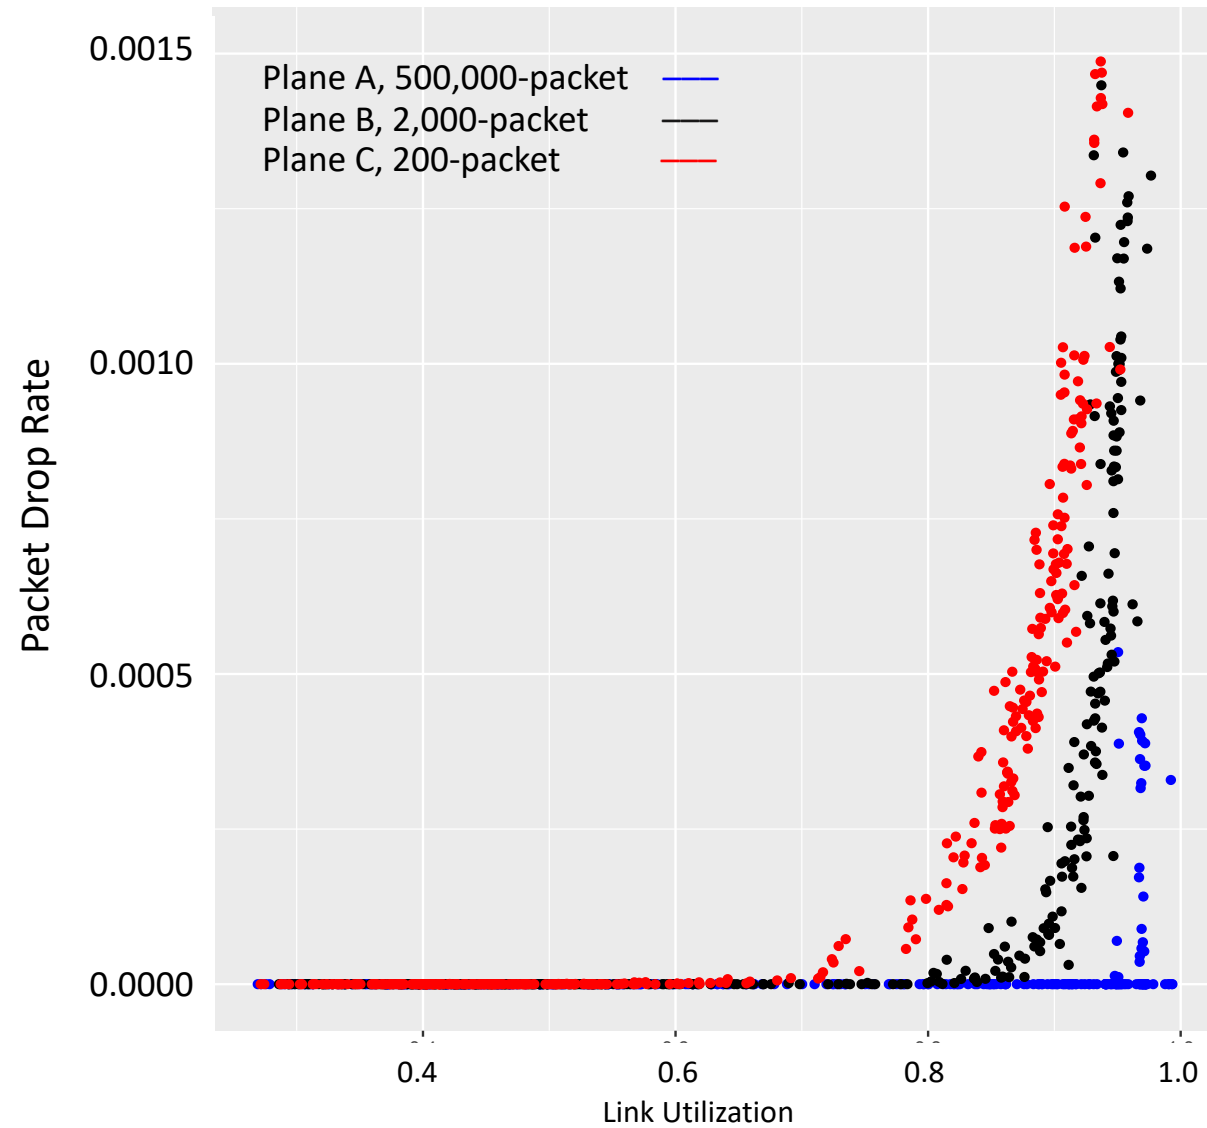


# RTT






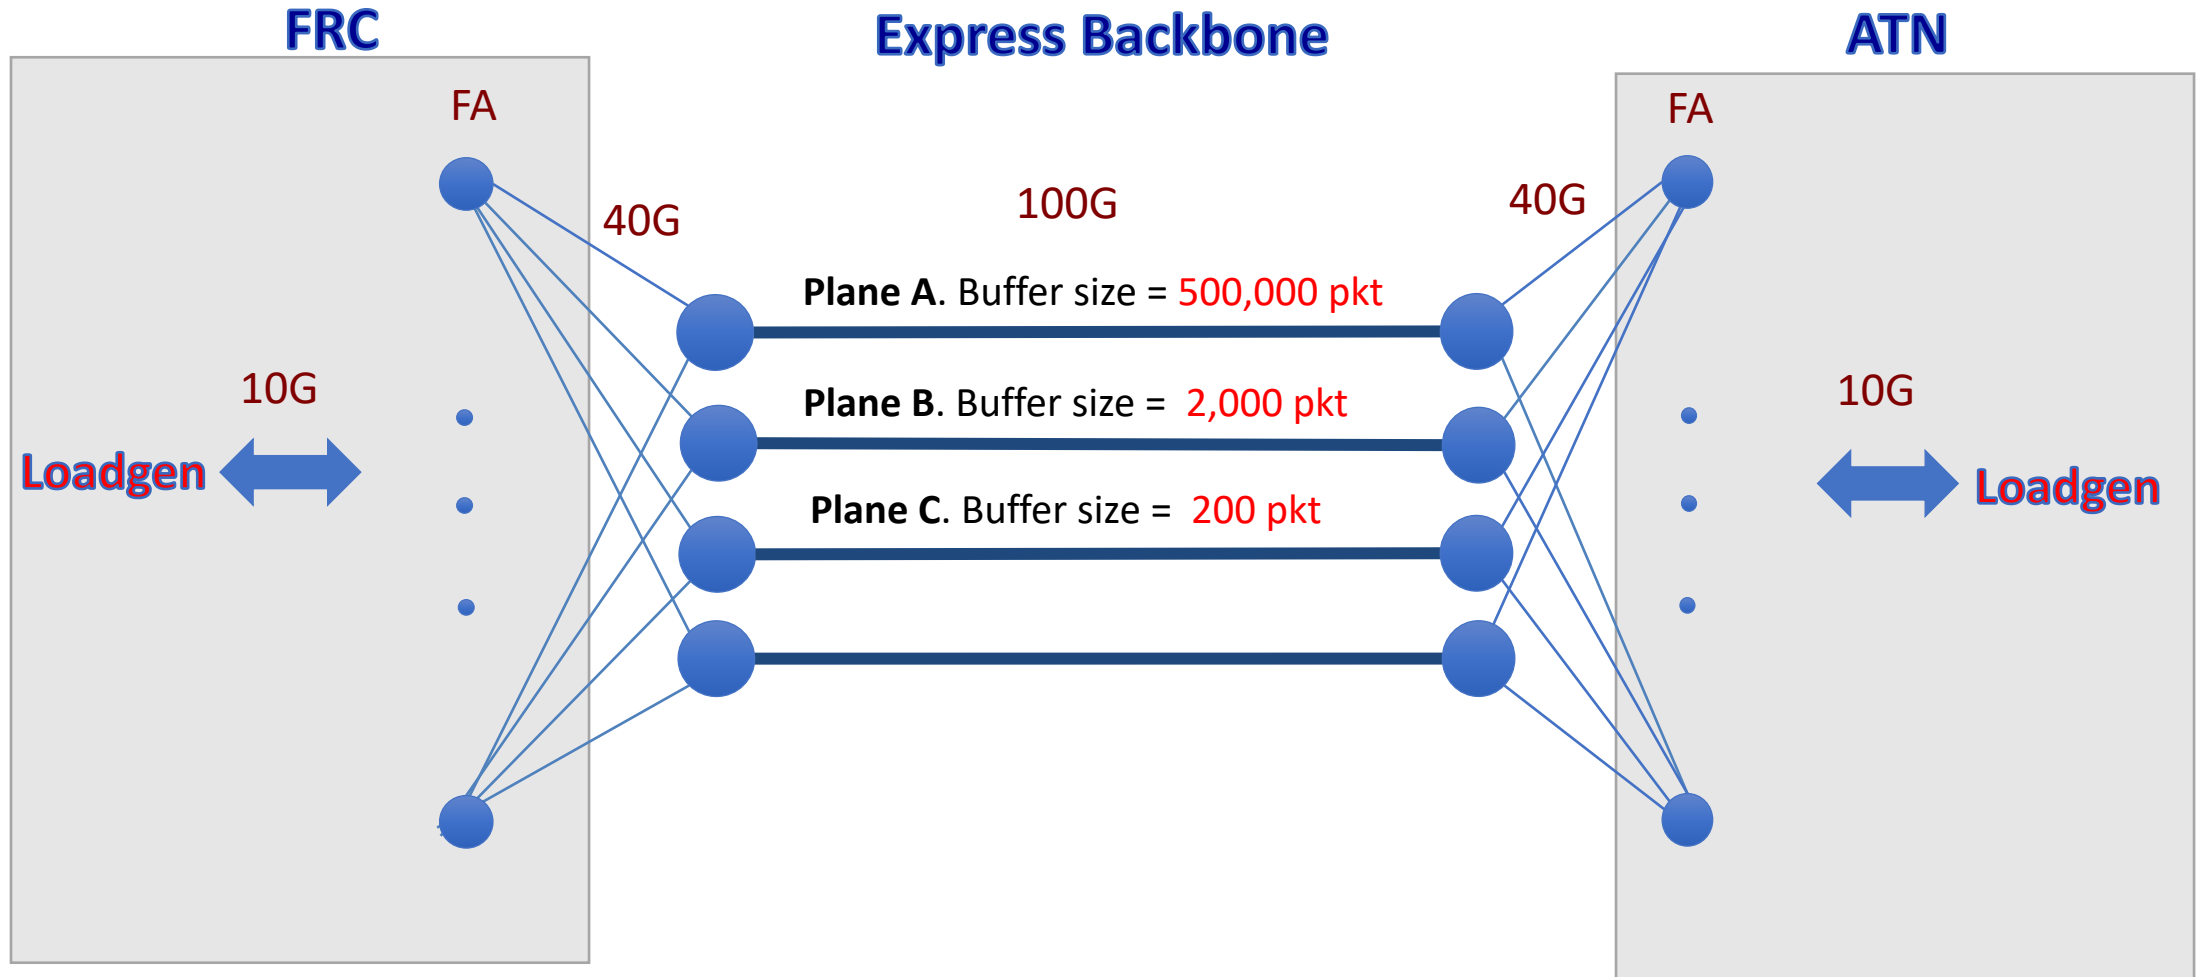
# Q1. Packet drops



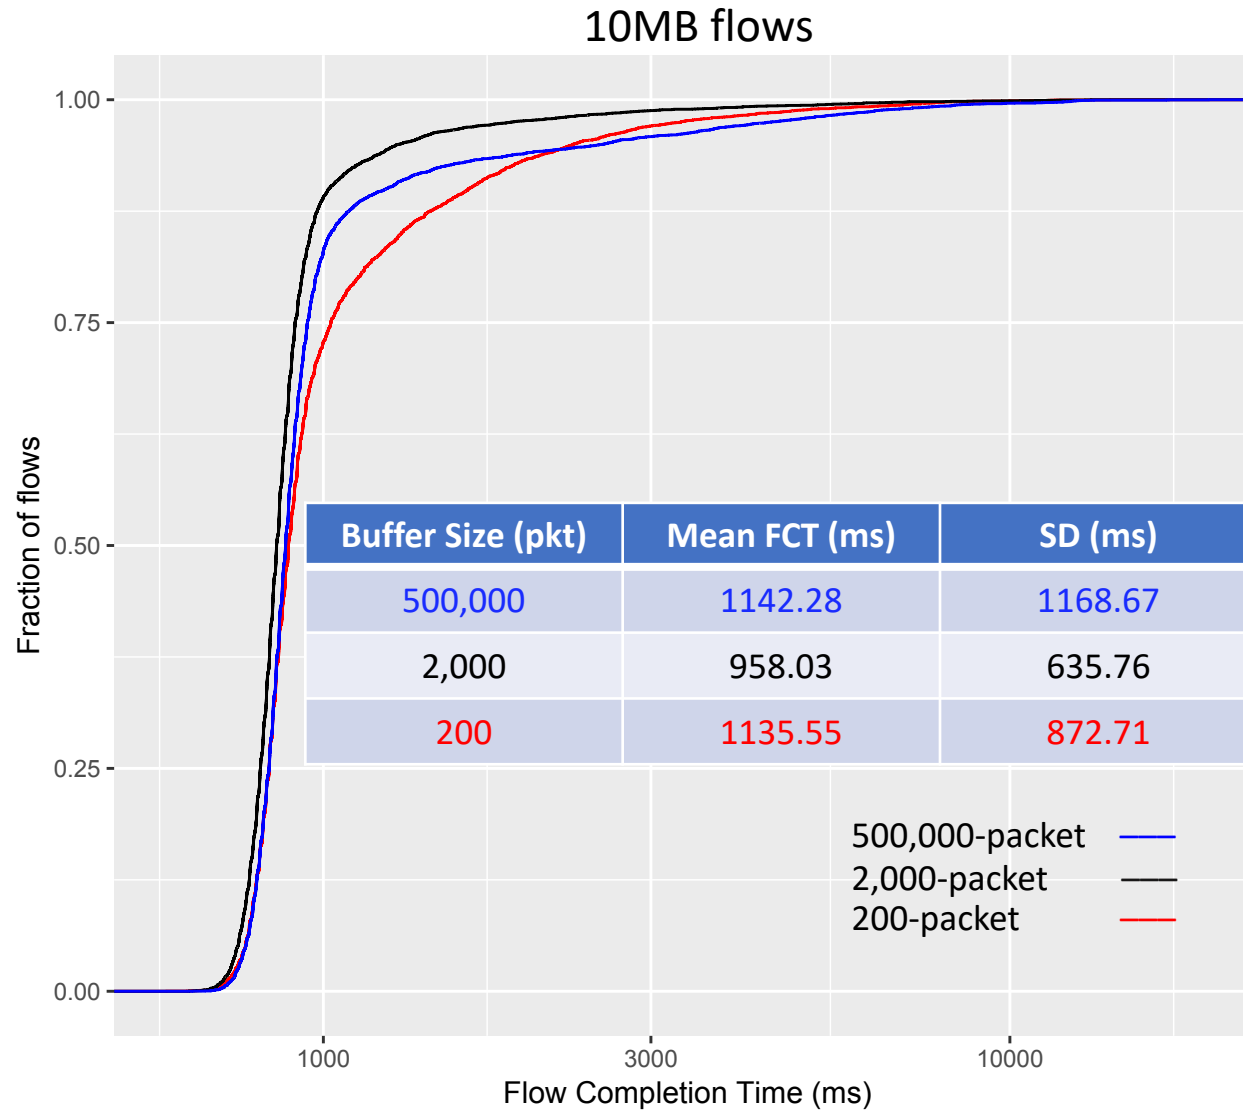
## Q2. Link utilization



# Experiments in the Backbone



# Q3. Individual flows



# Next steps

- Better understanding of buffer size impact on different applications
  - Throughput-bound, delay-sensitive, user-facing, machine-to-machine
- Using fine-grained packet capturing tools for better analysis of traffic
  - packet size distributions
  - packet pace statistics
  - flow and TCP-session distributions
- Exploring network-assisted congestion control mechanisms using programmable ASICs
- Exploring congestion control mechanisms such as BBR

# Questions?