


Router Buffer Sizes in the WAN

draft-ksubram-lmap-router-buffer-sizes-01

Kamala Subramaniam

Microsoft

IETF 93 Prague



I E T F[®]

PROBLEM

WHAT?

- Quantify Router Buffer Sizes in the WAN
- Link latencies ~40 to 150 milliseconds

WHY?

- Drive down \$/GB

HOW?

- Mine, and analyse empirical data

PREVIOUS WORK

Rule of Thumb

- Buffer Size = $RTT * C$
- RTT is the Round Trip Time
- C is the capacity of bottleneck link
- Holds true for a single TCP flow OR few synchronized TCP flows

Appenzeller

- Buffer Size = $(RTT * C) / \text{sqrt}(N)$
- N is the number of concurrent flows in a link
- Works for a few hundred unsynchronized flows

OBSERVATIONS FROM PREVIOUS WORK

Theoretical in nature

Holds true in simulated testbeds

Does not hold true in links running on a providers backbone

- Trans-pacific and trans-atlantic links of latencies of 150 and 90 ms with link utilization of < 30% show packet discards with small buffers
- WAN links within NA with large buffers and link utilization of 60—70% show packet discards

Need for new work

MAIN ISSUE

Lack of a standardized way to mine empirical data

Lack of a concise method to present mined data in a readable fashion

Data required for study of Router Buffer Sizes

Number of Concurrent Flows, N

Length of the Flow, L

Packet Discards, P

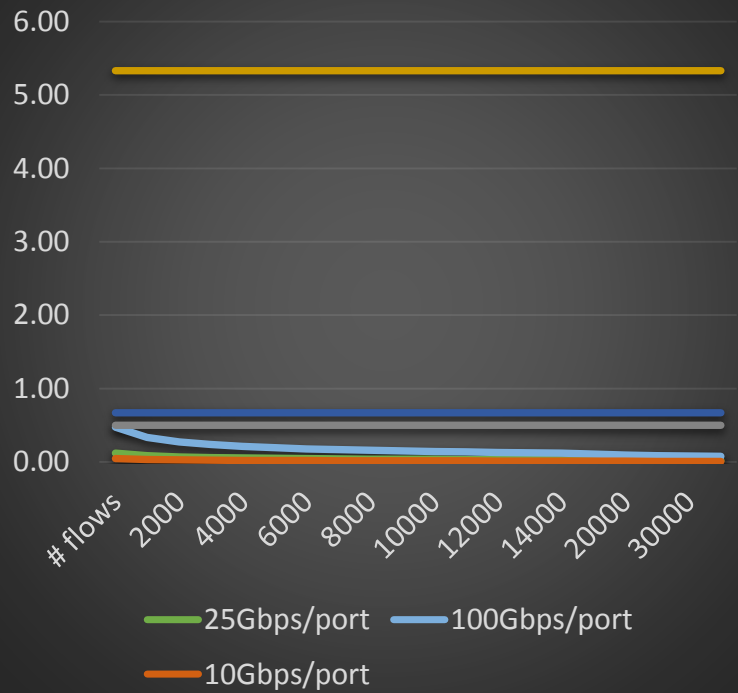
Reason for Packet Discards, R

Resolution of Time Interval, T

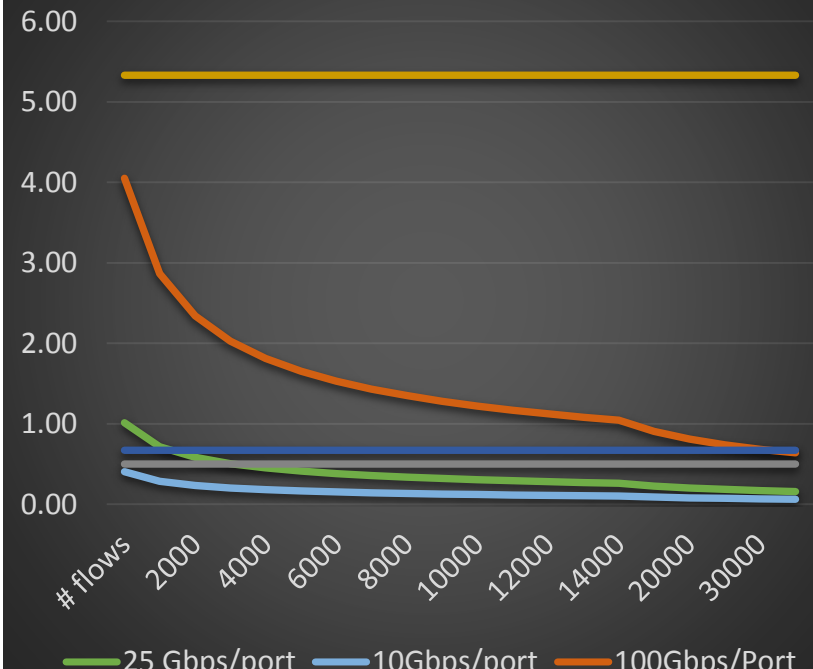
5 Tuple Flow Identity, I

Empirical Data

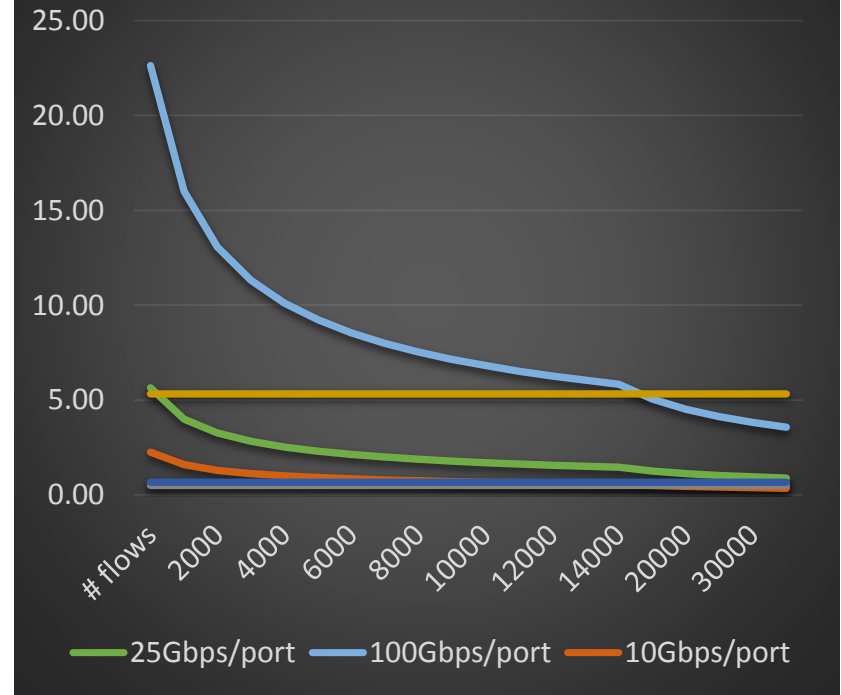
10ms RTT



86ms RTT



60ms RTT



Conclusion

Study of router buffer sizes is important, and incomplete

Need to mine empirical data for said study is important

Need to standardize these methods would be useful