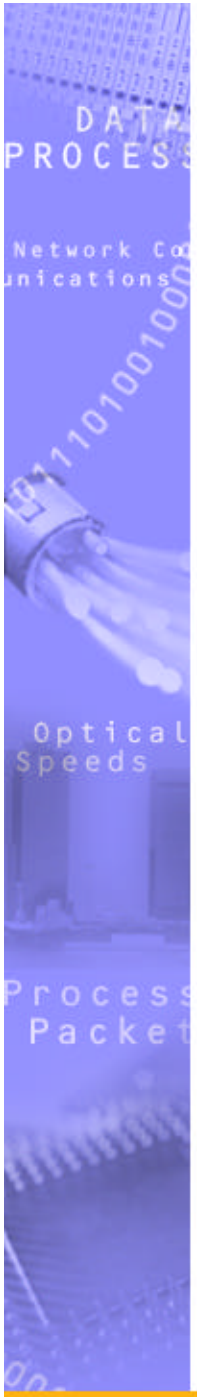

AuroraNetics, Inc.

Transit Buffer Implementations for 3 Classes of Service in RPR

Necdet Uzun and Pinar Yilmaz

AuroraNetics, Inc.

July 8-13, 2001
Portland, Oregon
802-17-01-00020, nu_crate_03.pdf

Contents

AuroraNetics, Inc.

- Objectives
- Why 3 priorities?
- Node Model
- Simulation results
- Conclusion

Objectives

AuroraNetics, Inc.

- Provide 3 priority classes in the ring:
- High Priority
  - Guaranteed bandwidth (provisioned)
  - Bounded delay and bounded jitter
- Medium Priority
  - Committed bandwidth (provisioned), best effort for excess traffic
  - Bounded delay and (loosely) bounded jitter
- Low Priority
  - No guarantees
  - Best effort for bandwidth, delay and jitter

Why 3 Priorities?

AuroraNetics, Inc.

- 3 traffic types with very specific behavior
- Mapped from Diffserv/MPLS/802.1Q CoS (3 bits)
- Good compromise between the types of services and the implementation
- Further classification performed by higher layers
- RPR MAC divides unprovisioned bandwidth fairly among nodes

Node Model

AuroraNetics, Inc.

- Mapped from Diffserv/MPLS/802.1Q (3 bits)

  If Priority < “Med Prio Threshold” then

  Low Priority

  Else if Priority < “High Prio Threshold” then

  Med Priority

  Else High Priority

- Committed Access Rate (CAR) for MP
  - MP Traffic only when exceeds CAR is subject to fairness algorithm control
- Number of transit buffers configurable
  - 2Tb: MP gets buffered in LTb
  - 3Tb: MP gets buffered in a separate queue

Node Model (cont)

AuroraNetics, Inc.

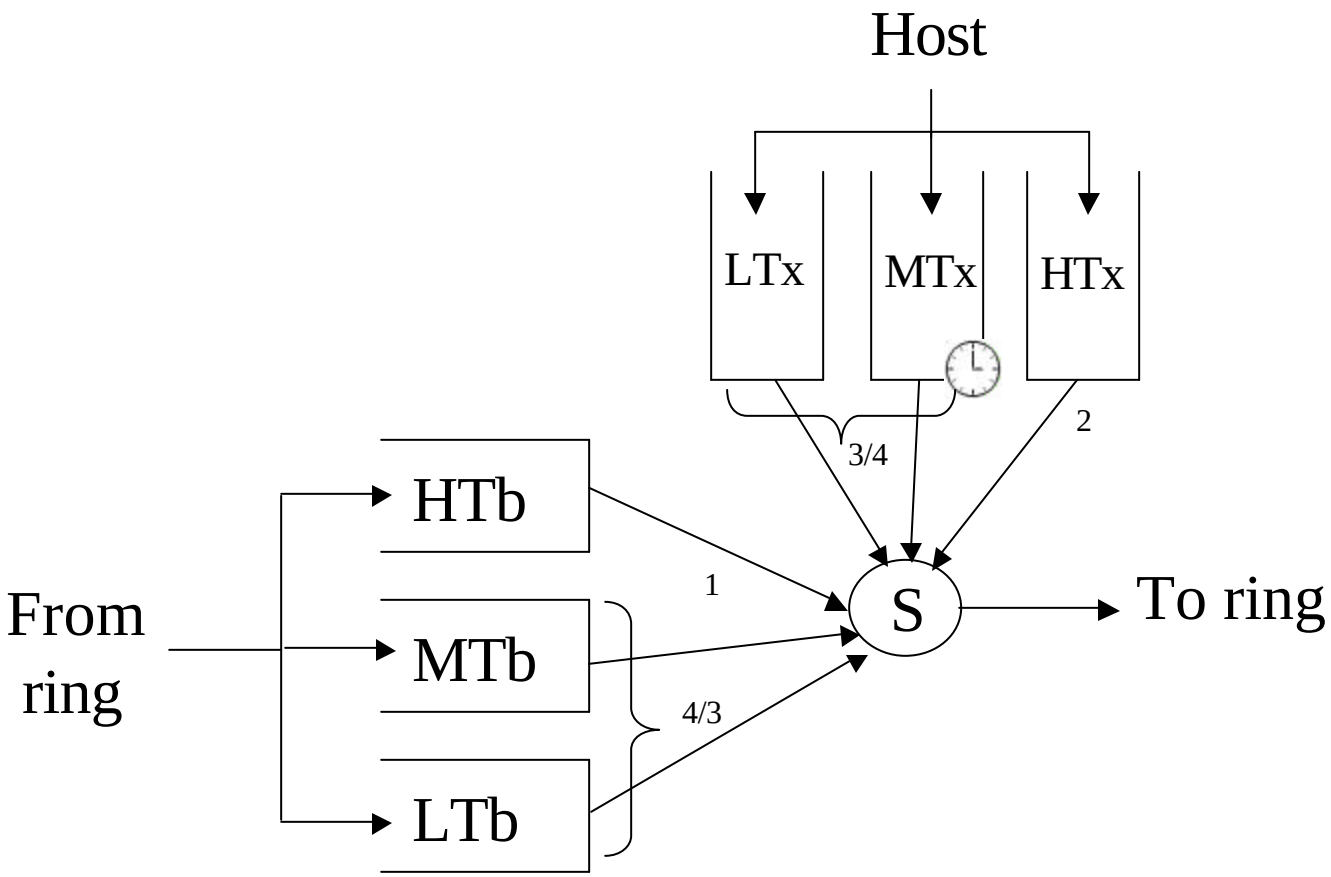

Important Statistics

AuroraNetics, Inc.

(5) (1) (2)
MAC ETE = Queuing Delay + Medium Access Delay +
Ring ETE Delay (+Receive Buffer Delay)
(3) (4)

Ring ETE = Pkt Tx + PropDelay + Transit Node Delay

Transit Node Delay = Pkt Handling Time +
(Insertion/Tb) Buffer Delay

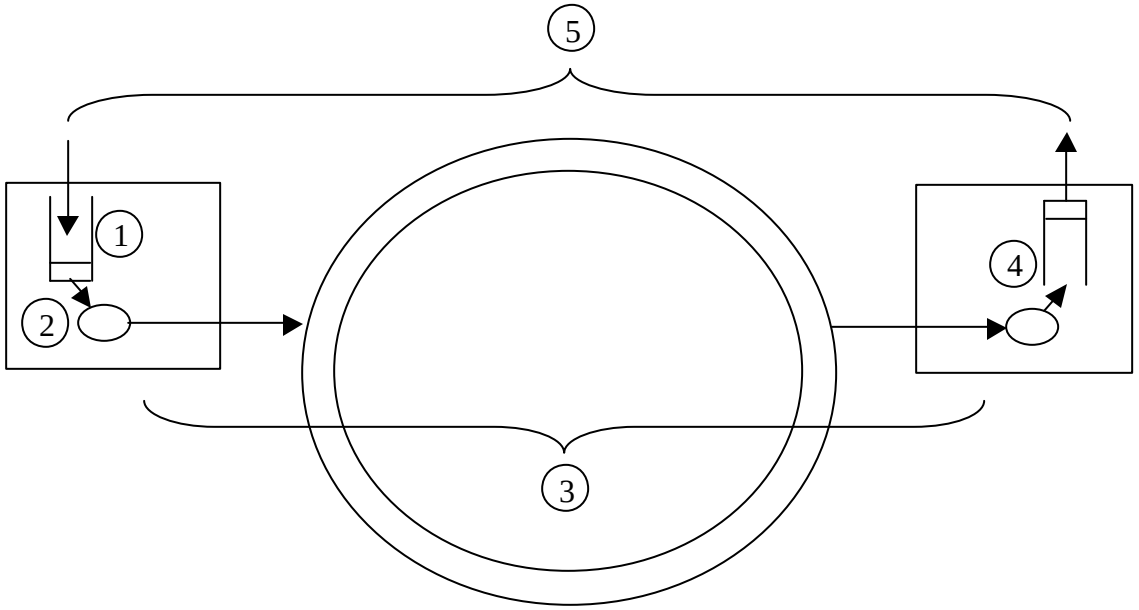

Scenarios

AuroraNetics, Inc.

- 16 node, 100Km OC192 dual ring
- Packet size: 64B(%60), 512B(%20), 1518B(%20)
- Buffers:
  - HTb = 32KB
  - MTb = 256KB
  - LTb = 256KB
  - LTx = 256MB
  - MTx = 512 KB
  - HTx = 256KB
  - Store and forward

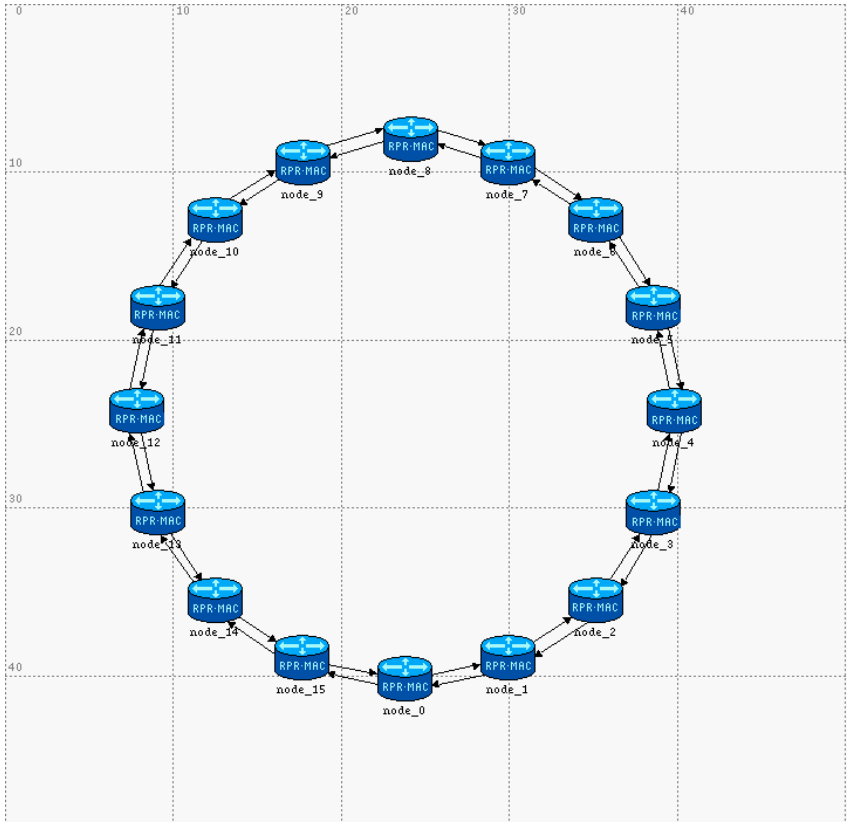

Hub Traffic Scenario

AuroraNetics, Inc.

- Node 7,6,5,4,3,2,1 to Hub (Node 0):
  - HTx: 430Mbps CBR
    - 1 tri-modal source per node
  - LTx: 1.15Gbps bursts
    - 1 tri-modal source per node, on 1msec, off 1msec, exponential distribution
    - total of ~570Mbps LTx per node
  - MTx: 860Mbps bursts
    - 1 tri-modal source per node, on 1msec, off 1msec, exponential distribution
    - total of ~430Mbps MTx per node
    - CAR per node 430Mbps
- In addition, Nodes double their MP traffic at 1sec
- Total traffic destined at Hub: ~10Gbps
  - 0.1-1 sec: %30HP, %30MP, %40LP
  - 1-2 sec: %30HP, %60MP, %10LP
  - Total simulation run-time: 2sec

Throughput Comparison

AuroraNetics, Inc.

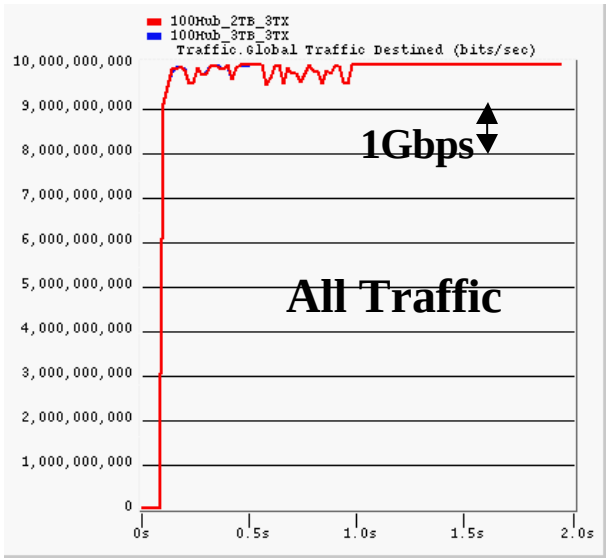

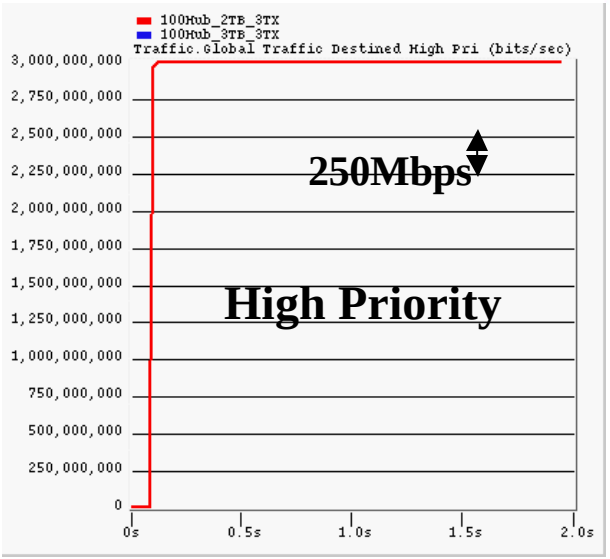

2TB

3TB

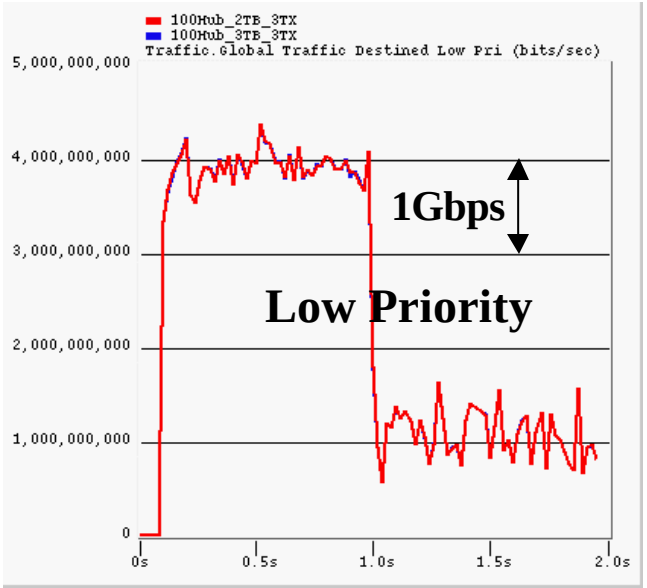

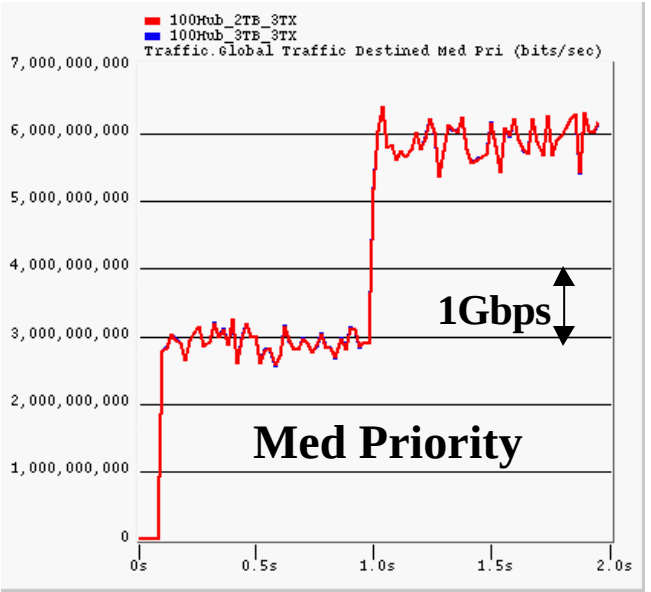

High Priority Delay Comparison

AuroraNetics, Inc.

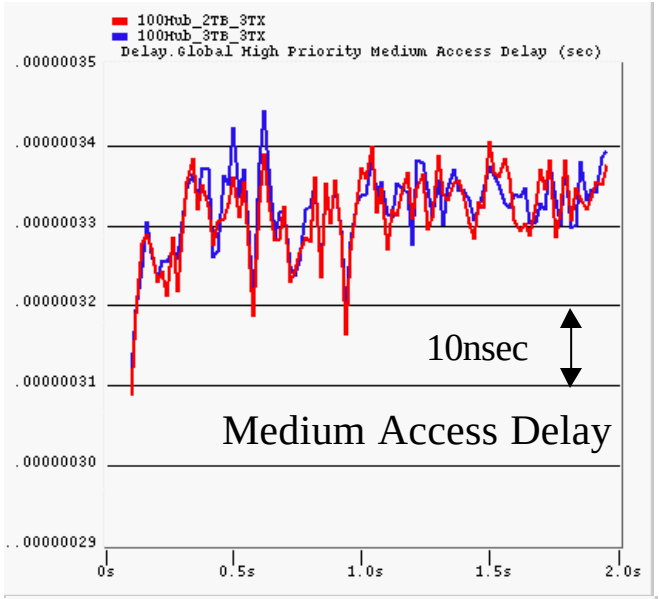

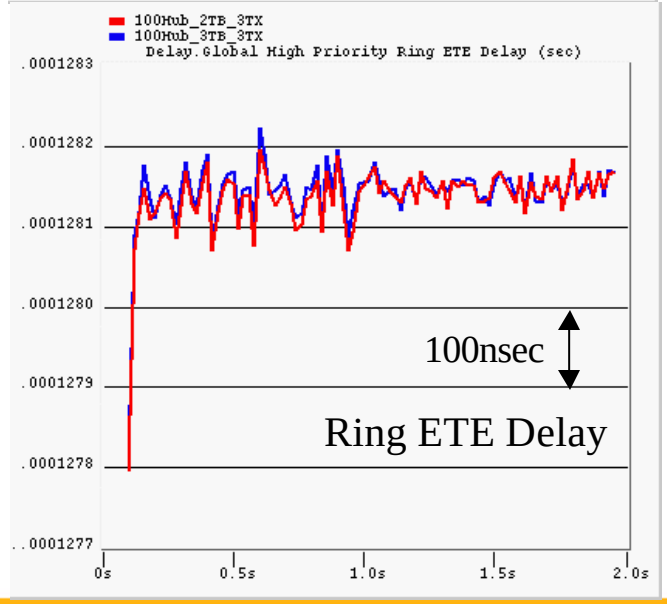

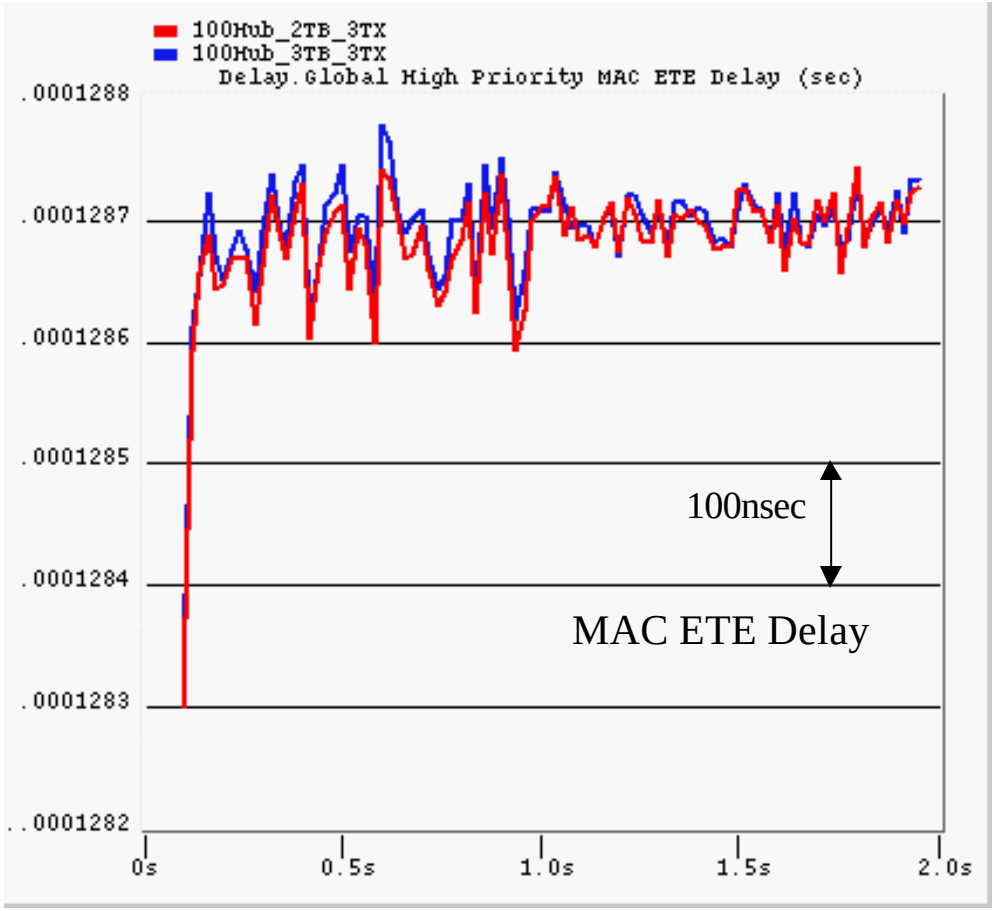

2TB

3TB

Med Priority Delay Comparison

AuroraNetics, Inc.

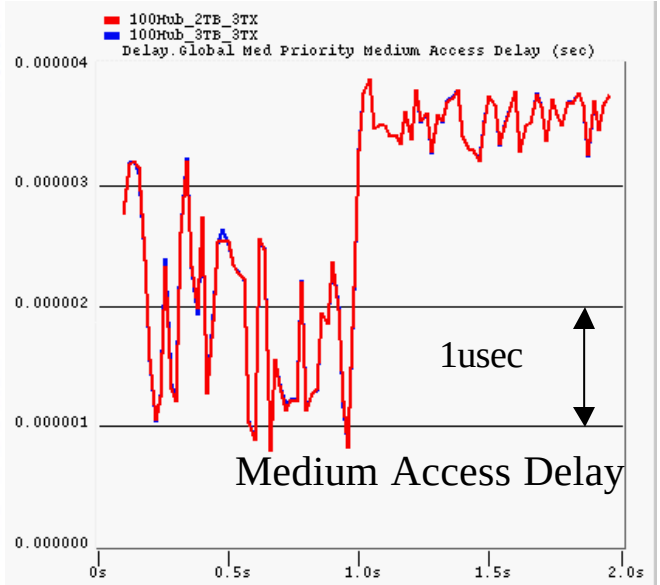

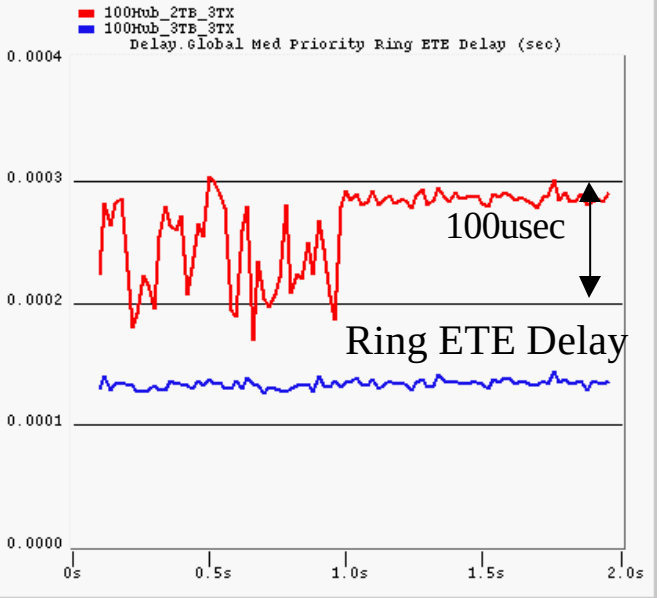

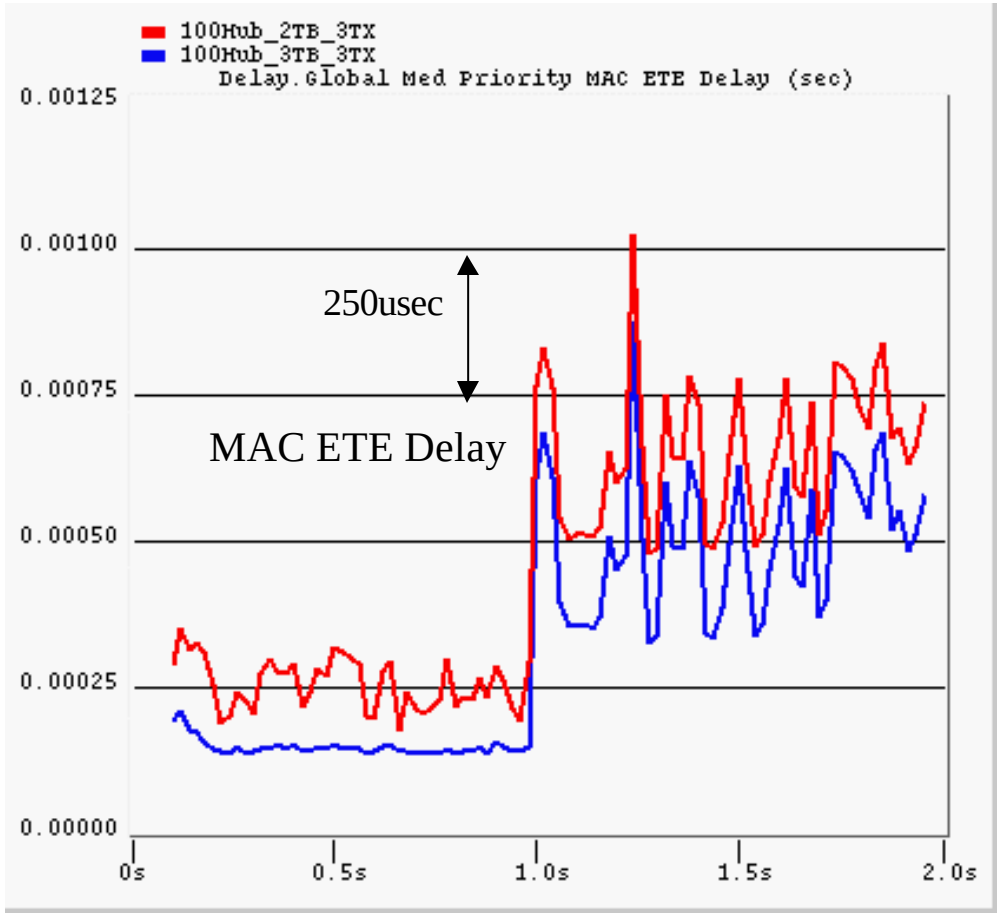

2TB

3TB

Low Priority Delay Comparison

AuroraNetics, Inc.

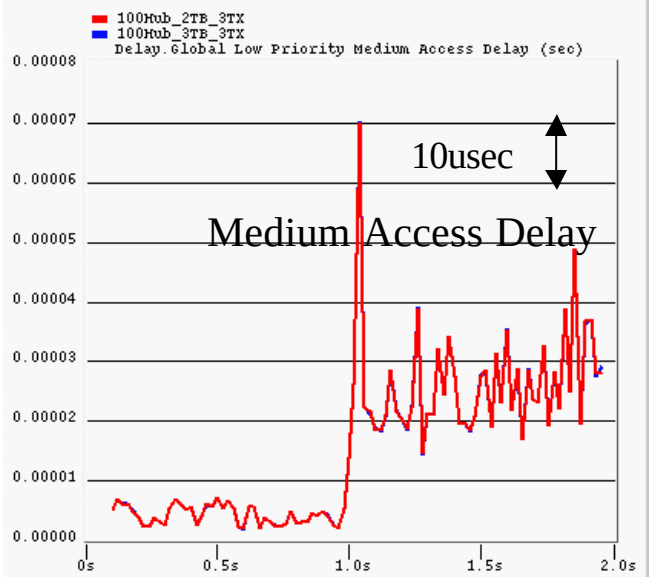

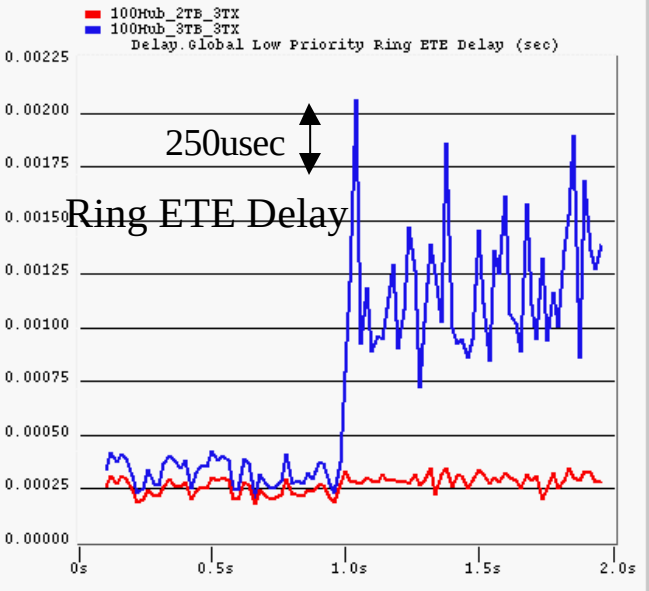

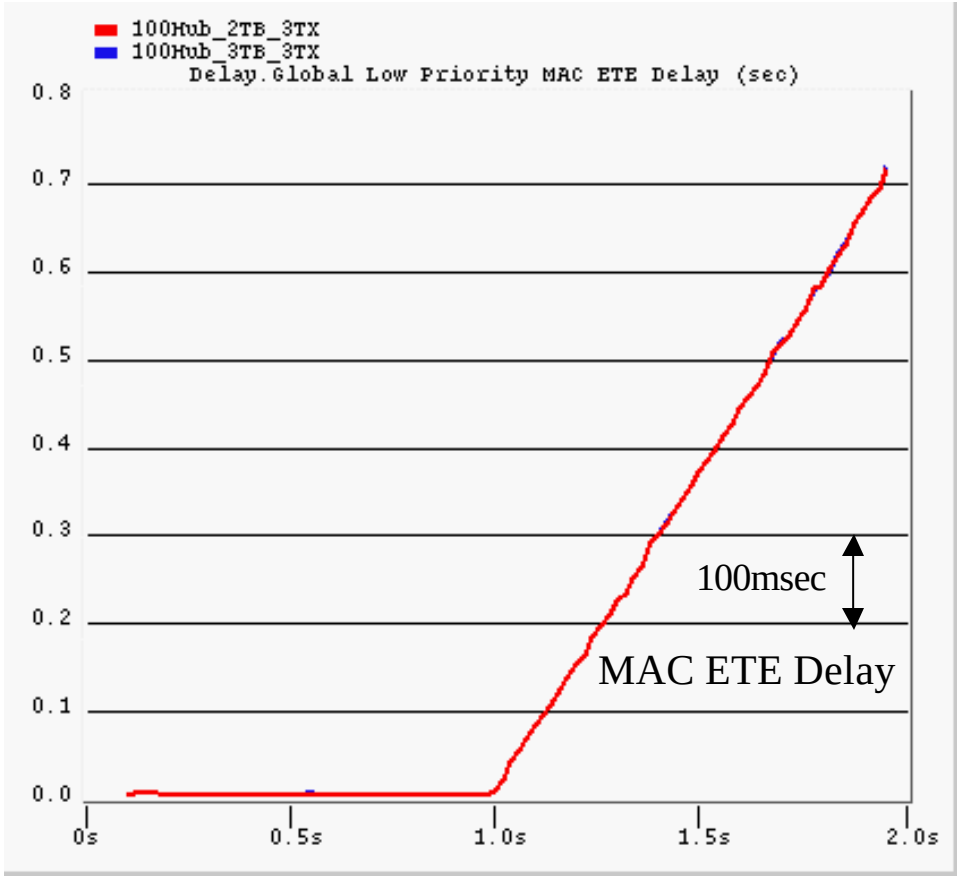

2TB

3TB

High Pri MAC ETE Delay Histogram

AuroraNetics, Inc.

3 Transit Buffers

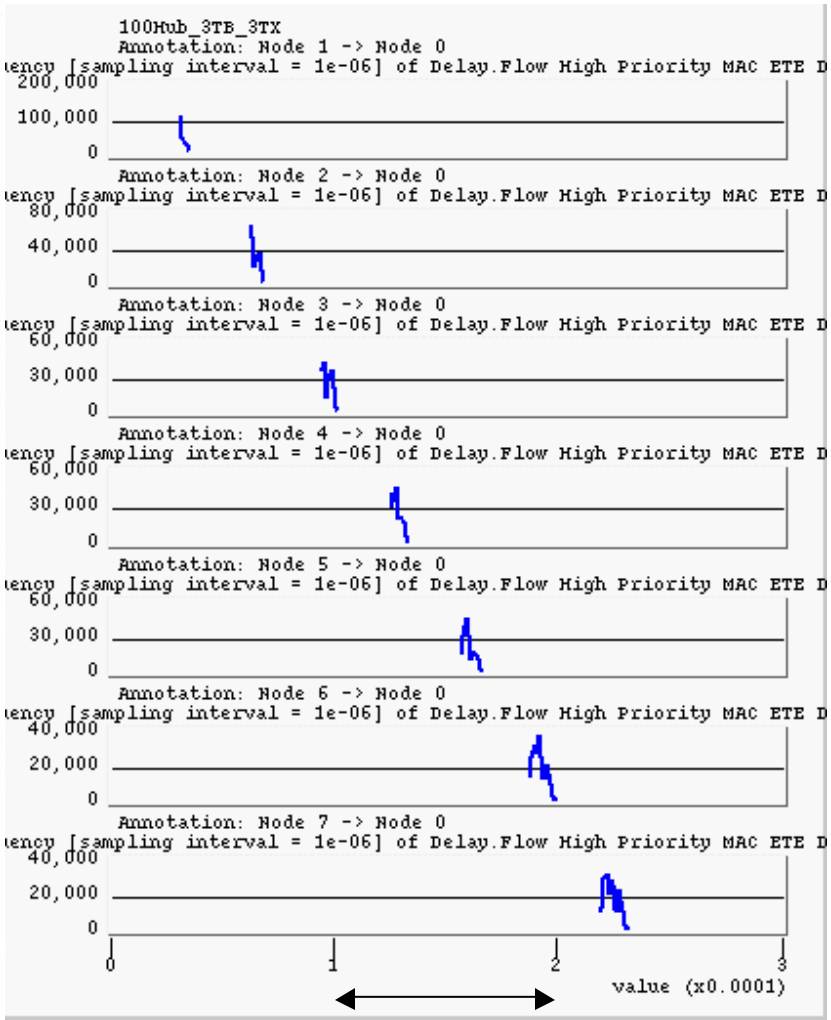

100usec

2 Transit Buffers

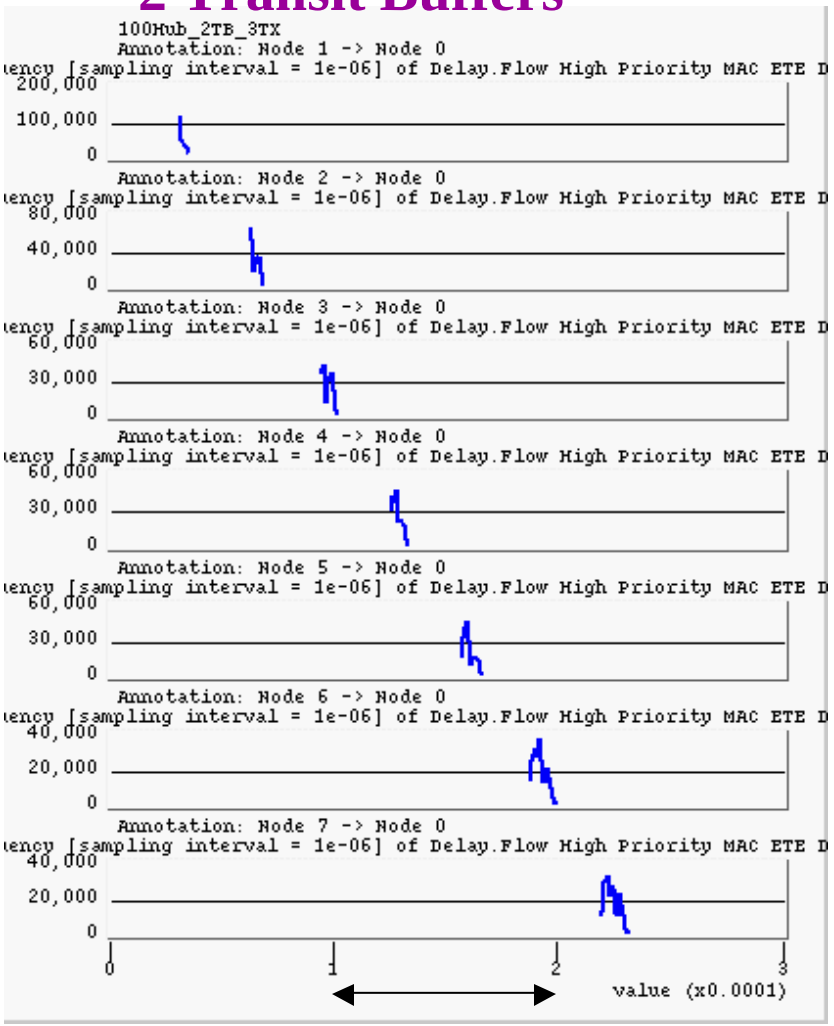

100usec

High Pri Ring ETE Delay Histogram

AuroraNetics, Inc.

3 Transit Buffers

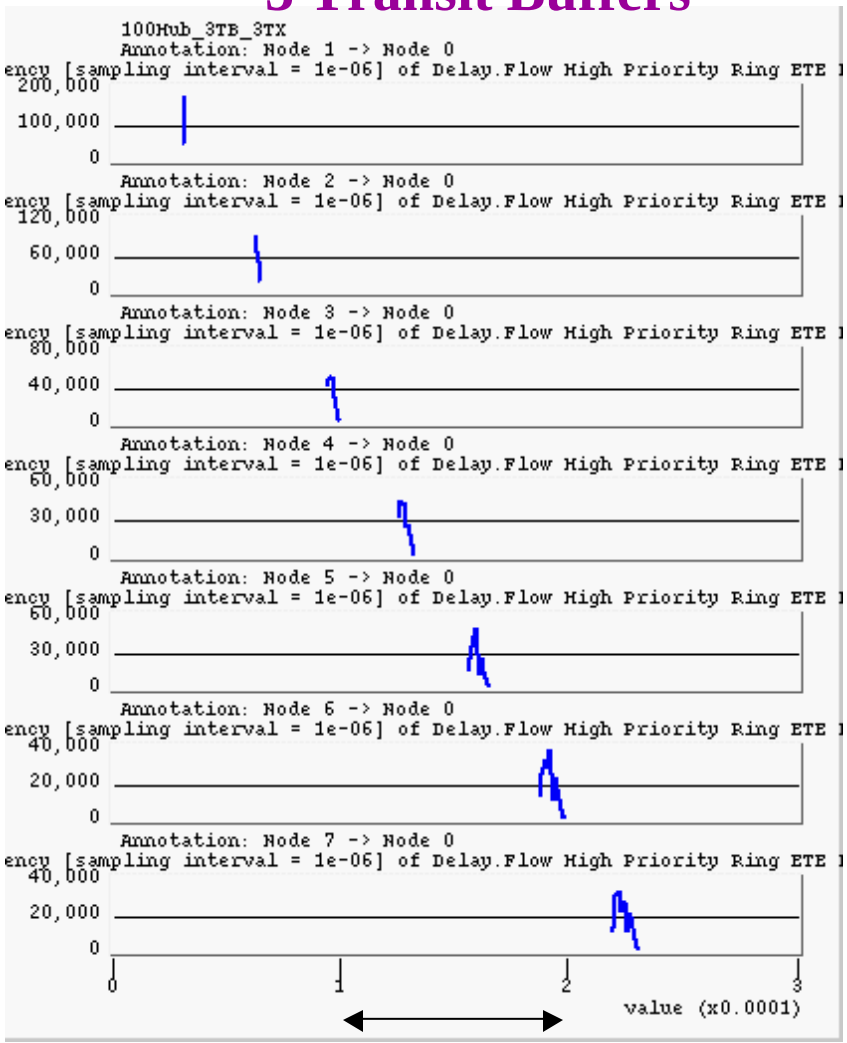

100usec

2 Transit Buffers

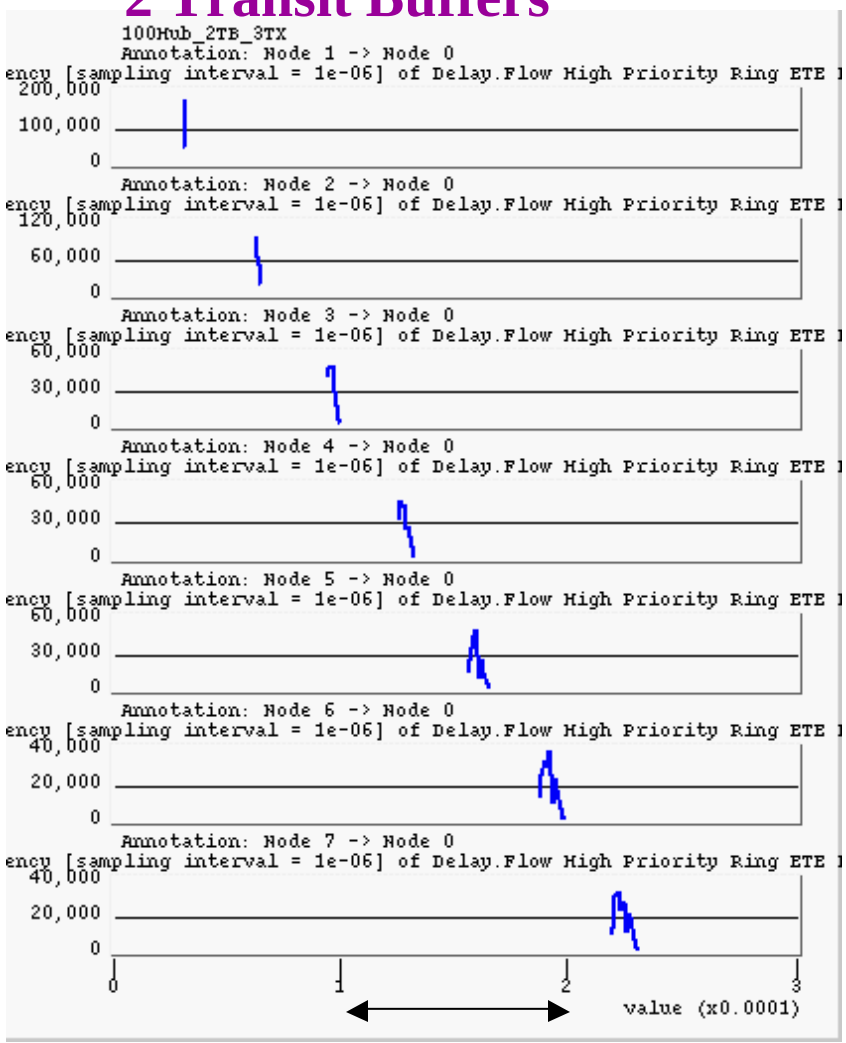

100usec

High Pri MedAcc Delay Histogram

AuroraNetics, Inc.

3 Transit Buffers

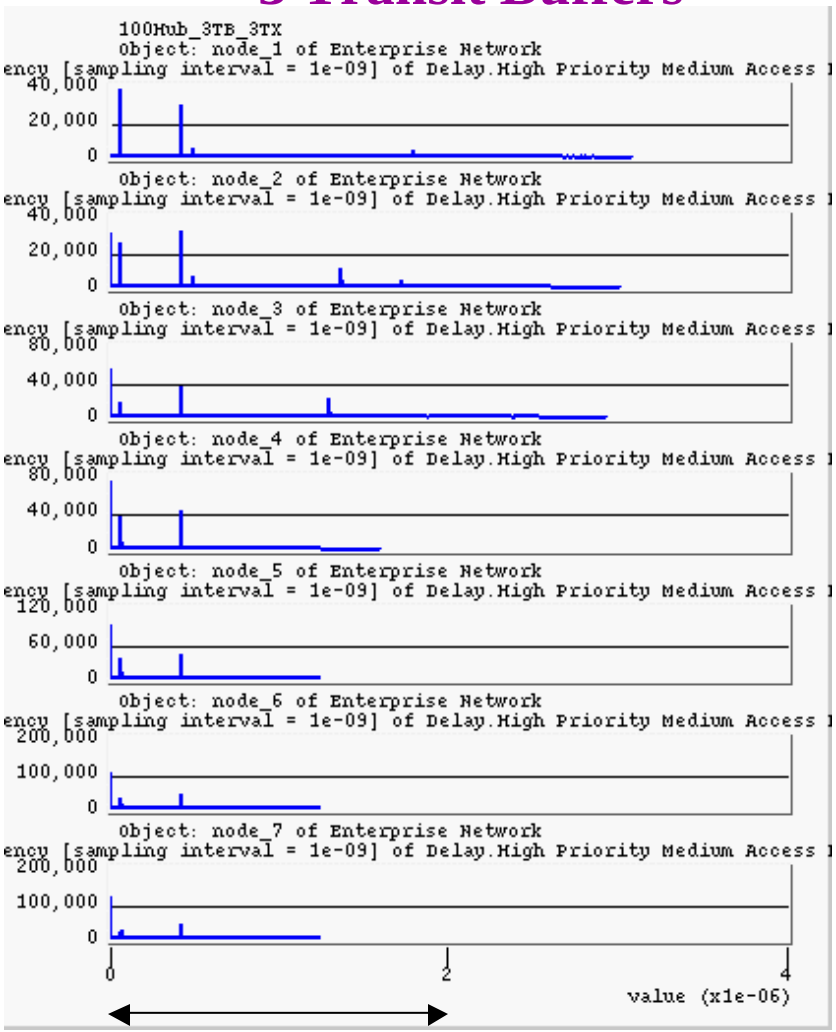

2usec

2 Transit Buffers

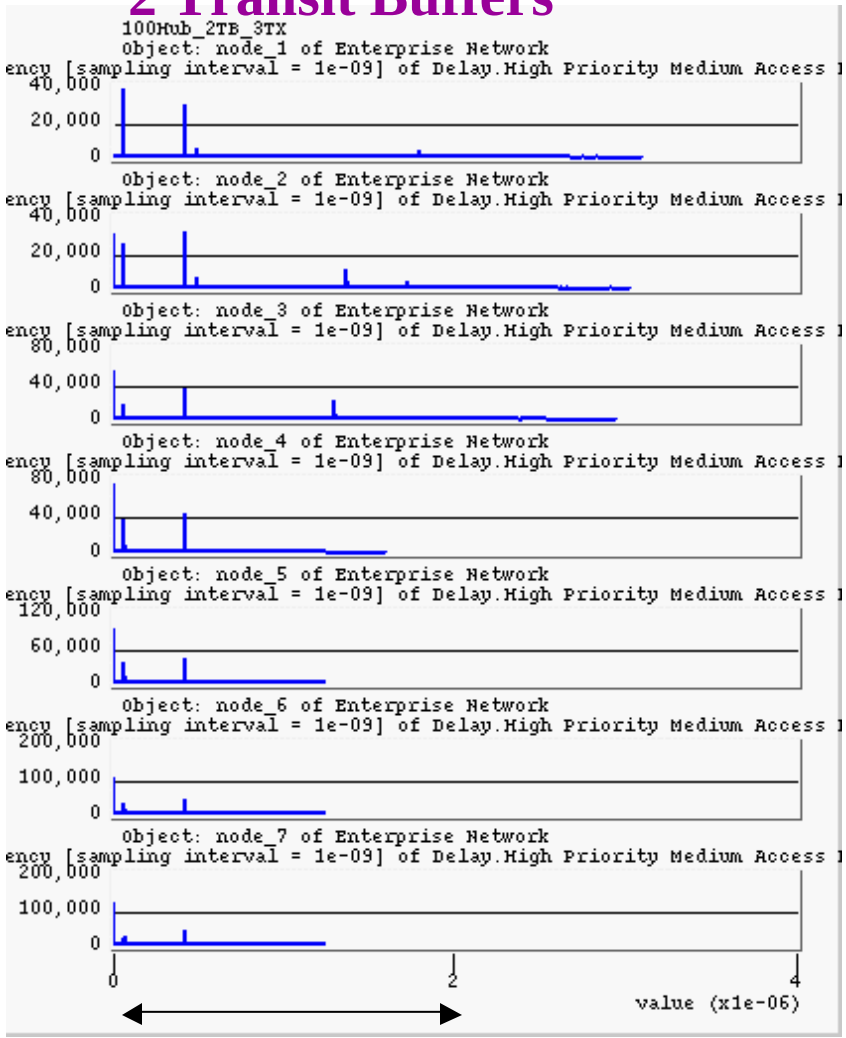

2usec

Med Pri MAC ETE Delay Histogram

AuroraNetics, Inc.

3 Transit Buffers

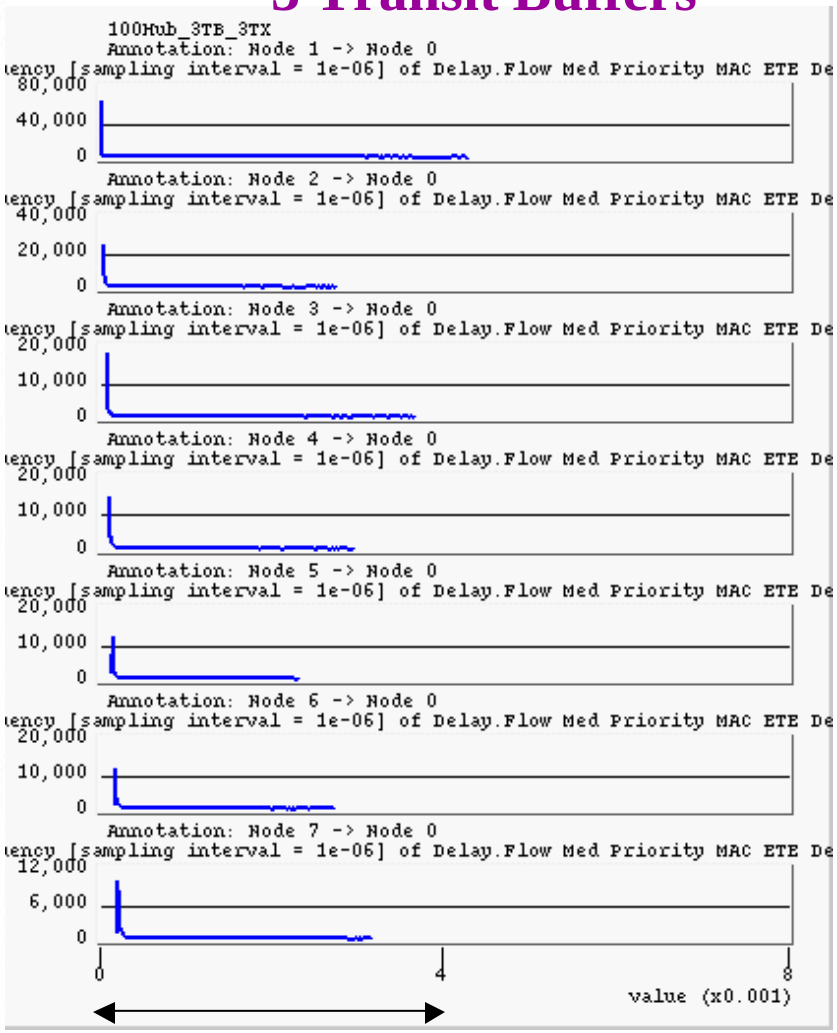

4msec

2 Transit Buffers

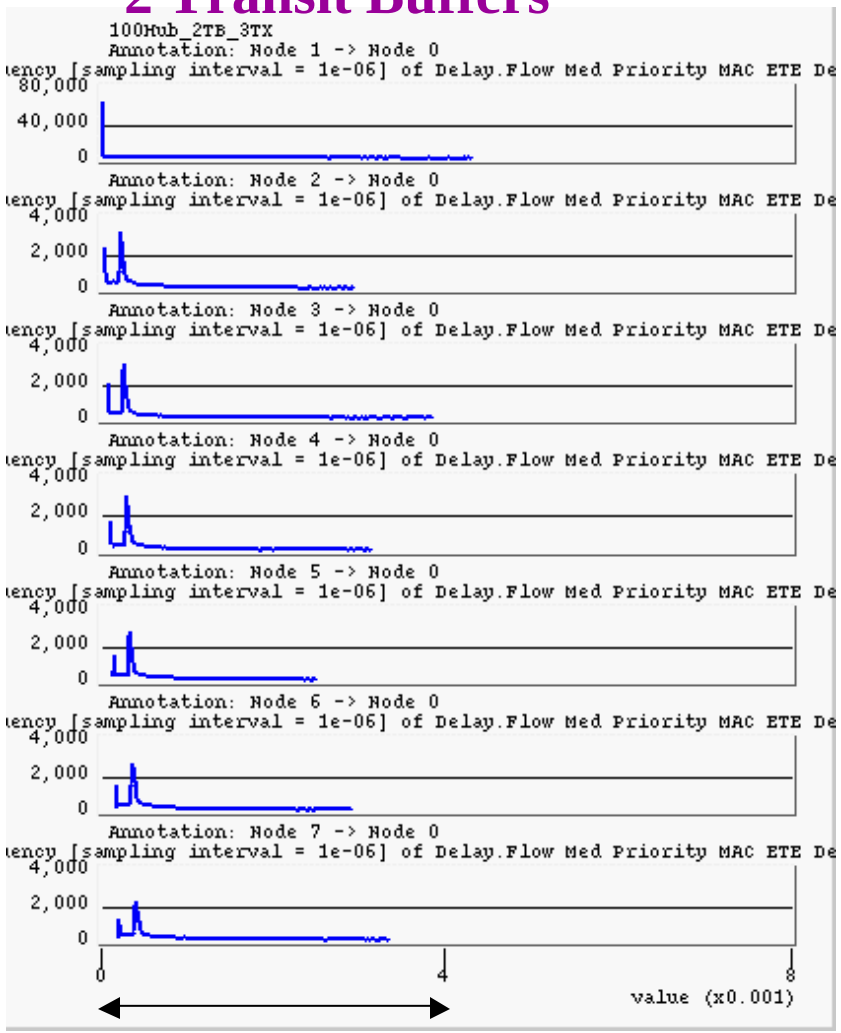

4msec

Med Pri Ring ETE Delay Histogram

AuroraNetics, Inc.

3 Transit Buffers

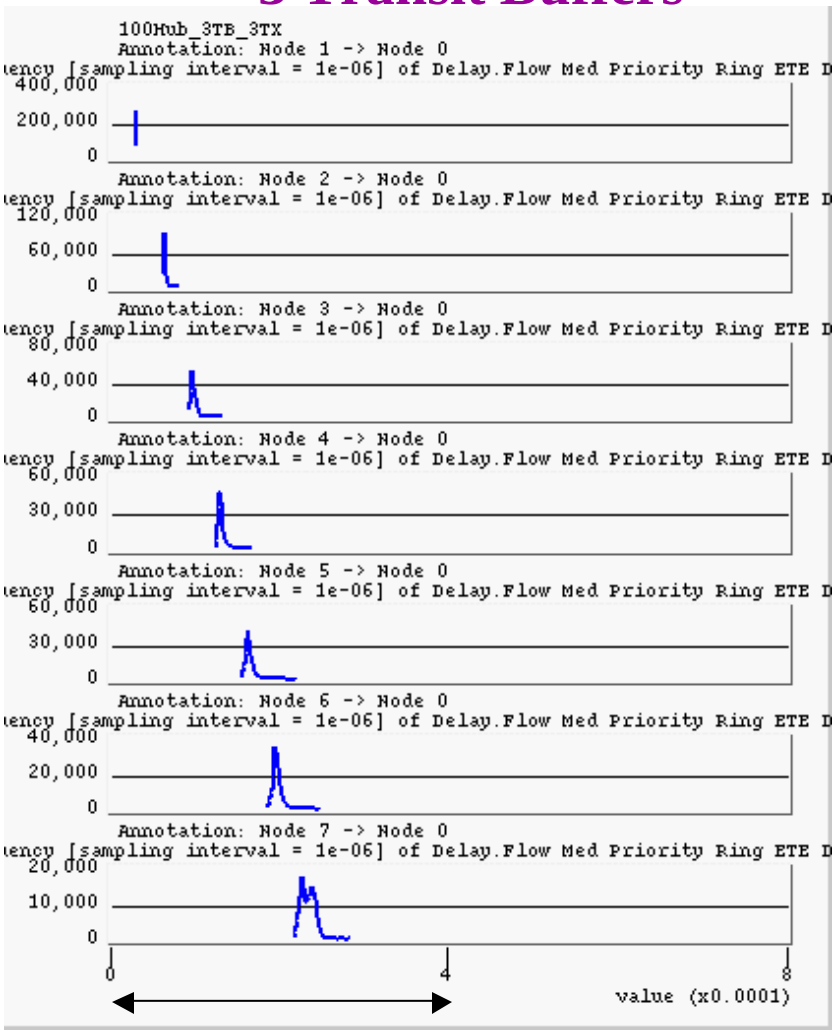

400usec

2 Transit Buffers

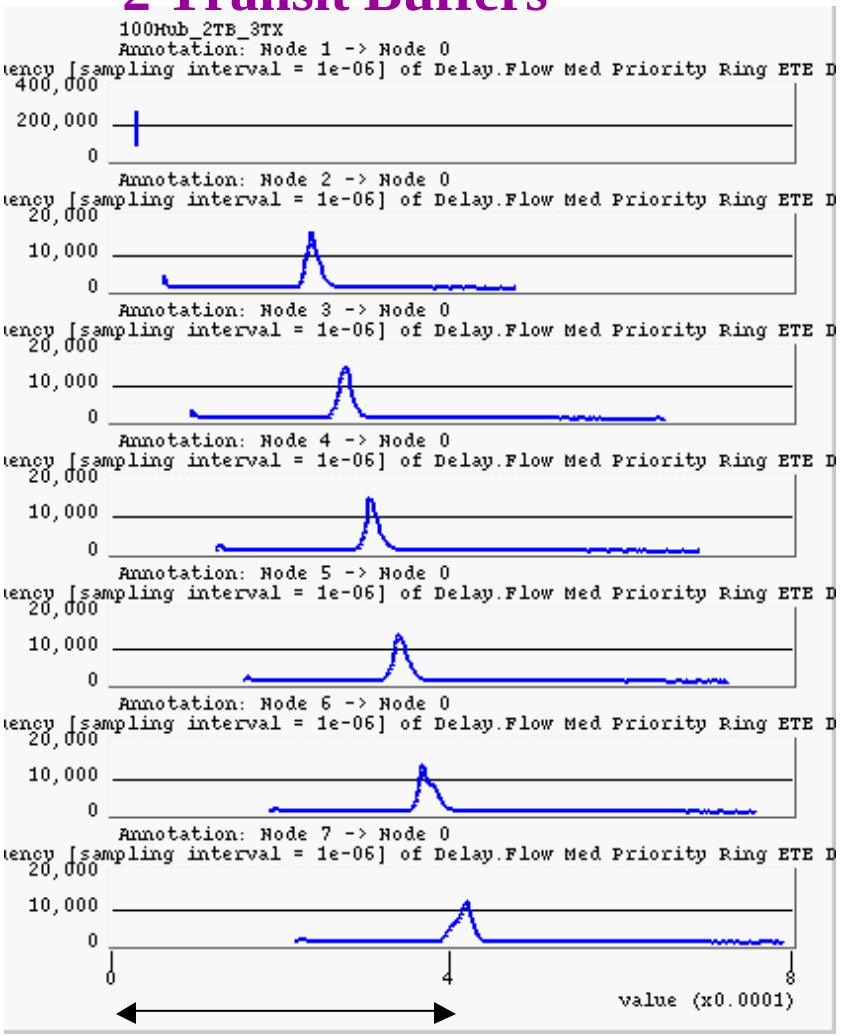

400usec

Med Pri MedAcc Delay Histogram

AuroraNetics, Inc.

3 Transit Buffers

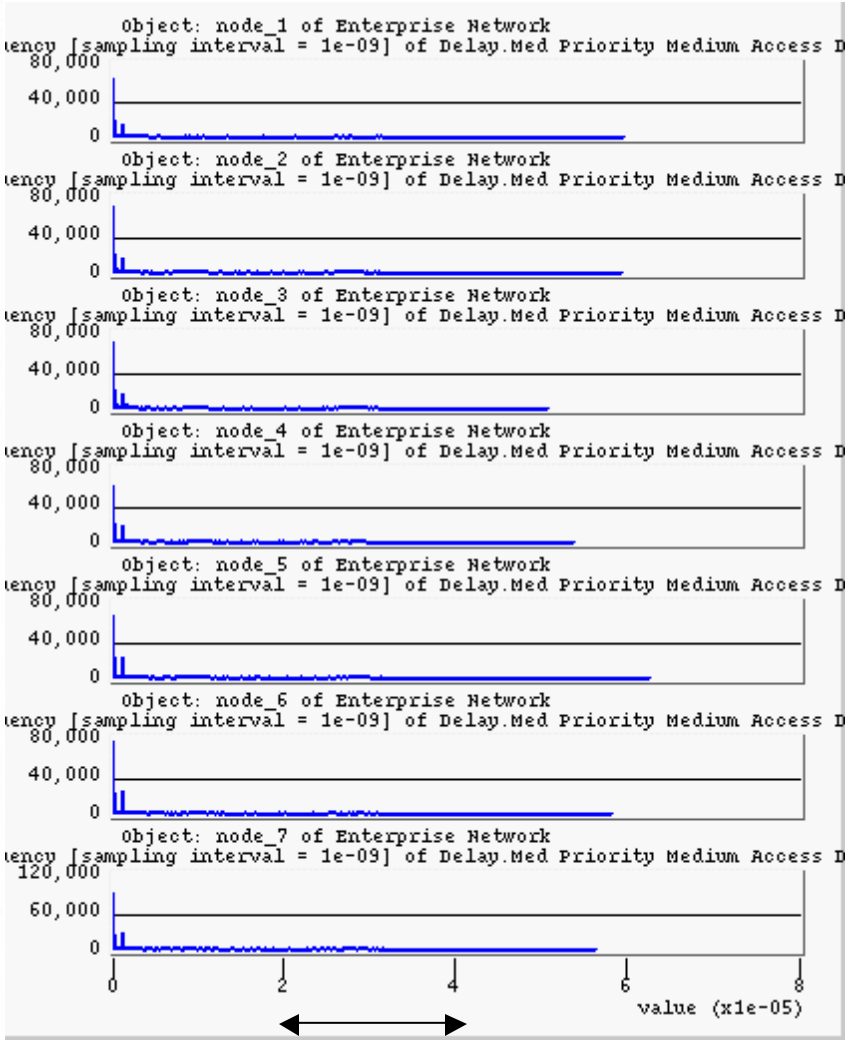

20usec

2 Transit Buffers

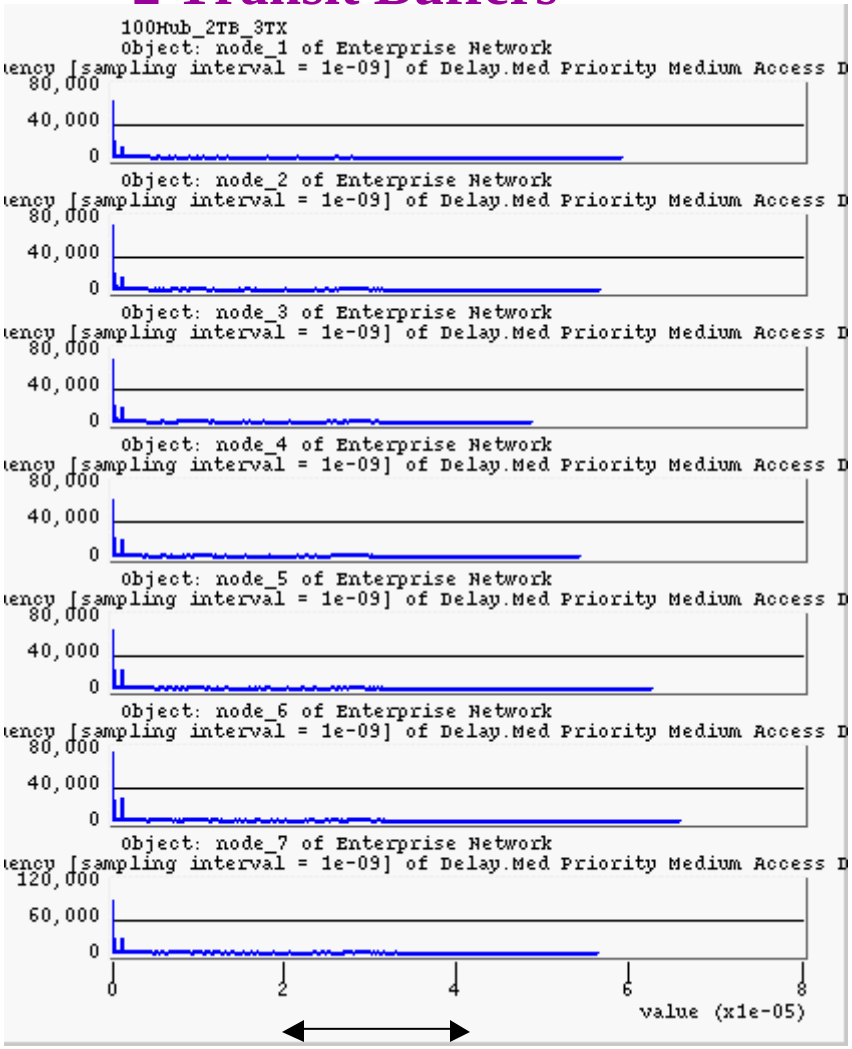

20usec

Low Pri MAC ETE Delay Histogram

AuroraNetics, Inc.

3 Transit Buffers

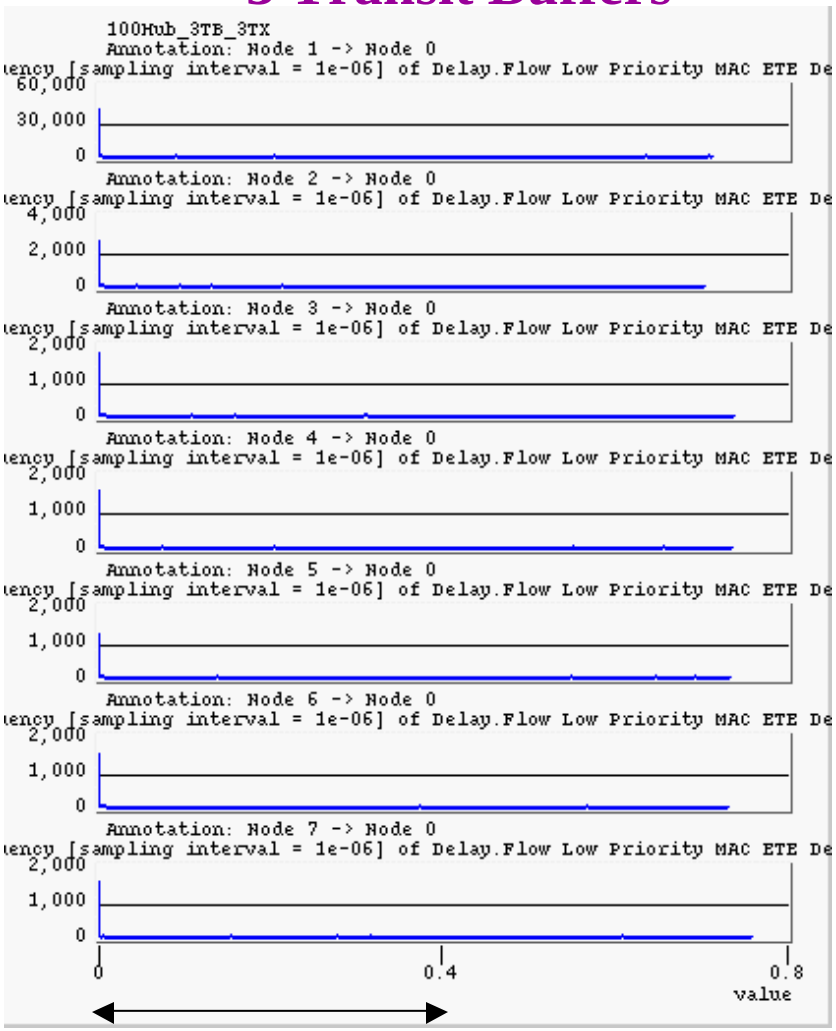

400msec

2 Transit Buffers

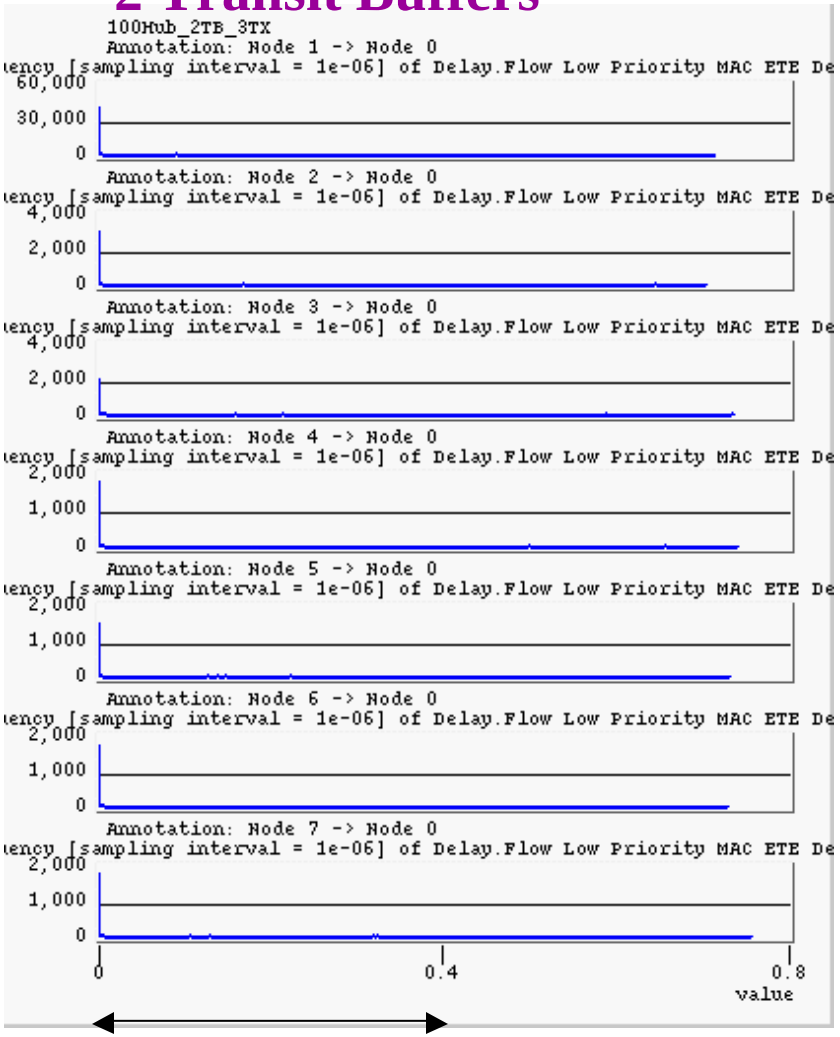

400msec

Low Pri Ring ETE Delay Histogram

AuroraNetics, Inc.

3 Transit Buffers

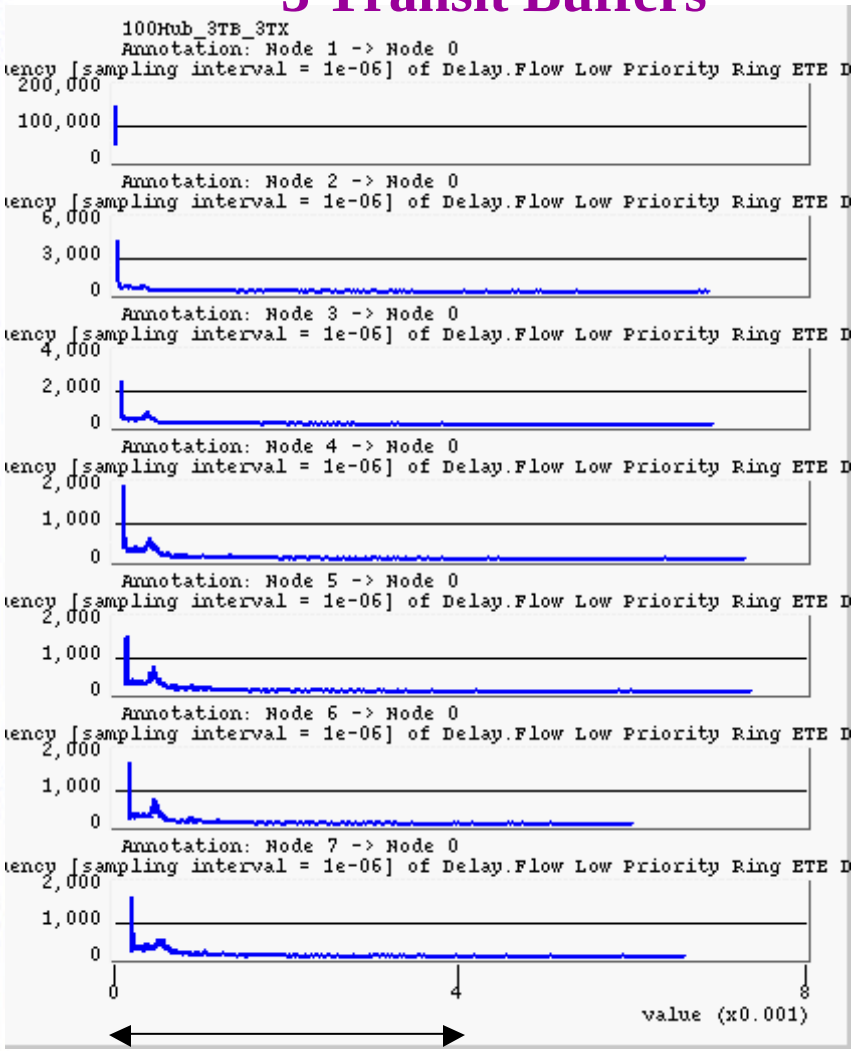

4msec

2 Transit Buffers

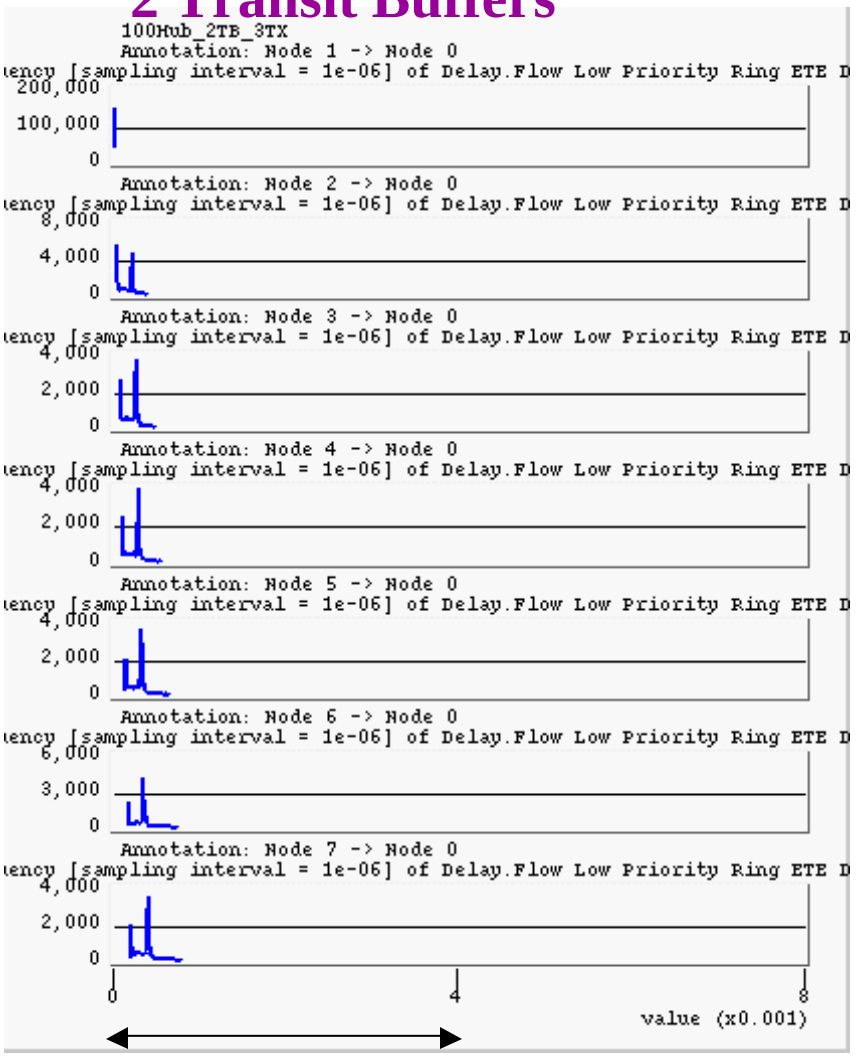

4msec

Low Pri MedAcc Delay Histogram

AuroraNetics, Inc.

3 Transit Buffers

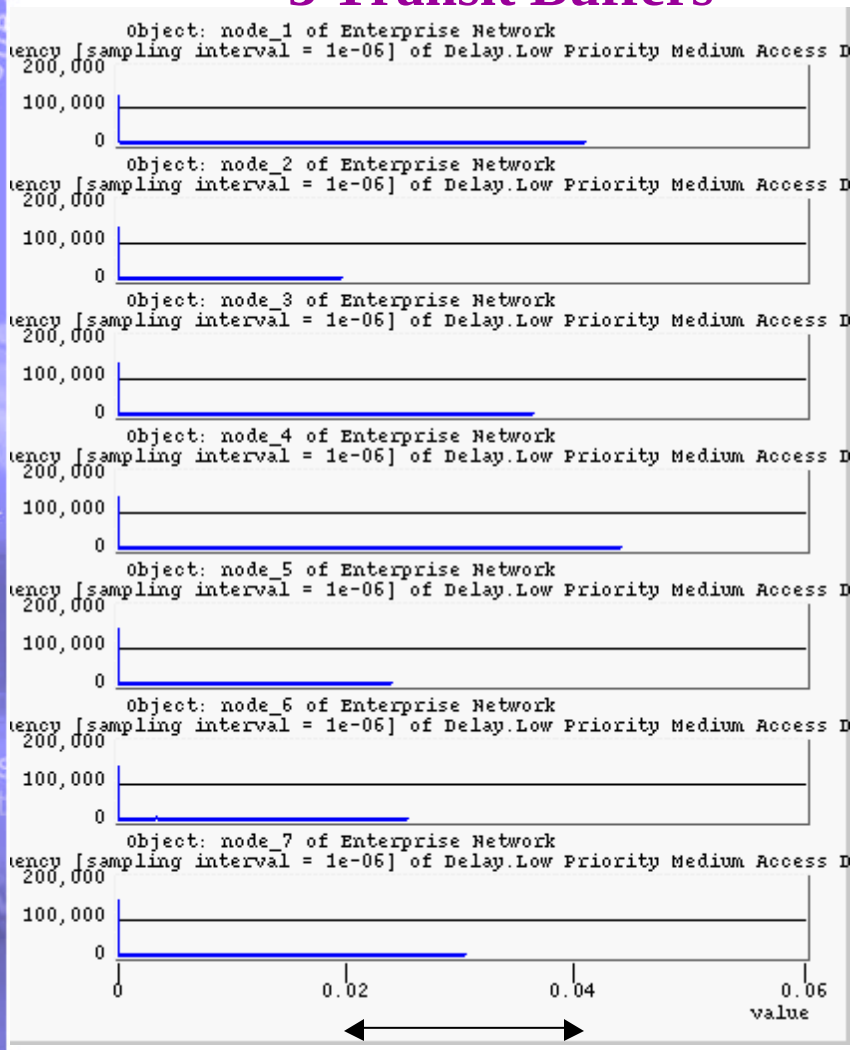

20msec

2 Transit Buffers

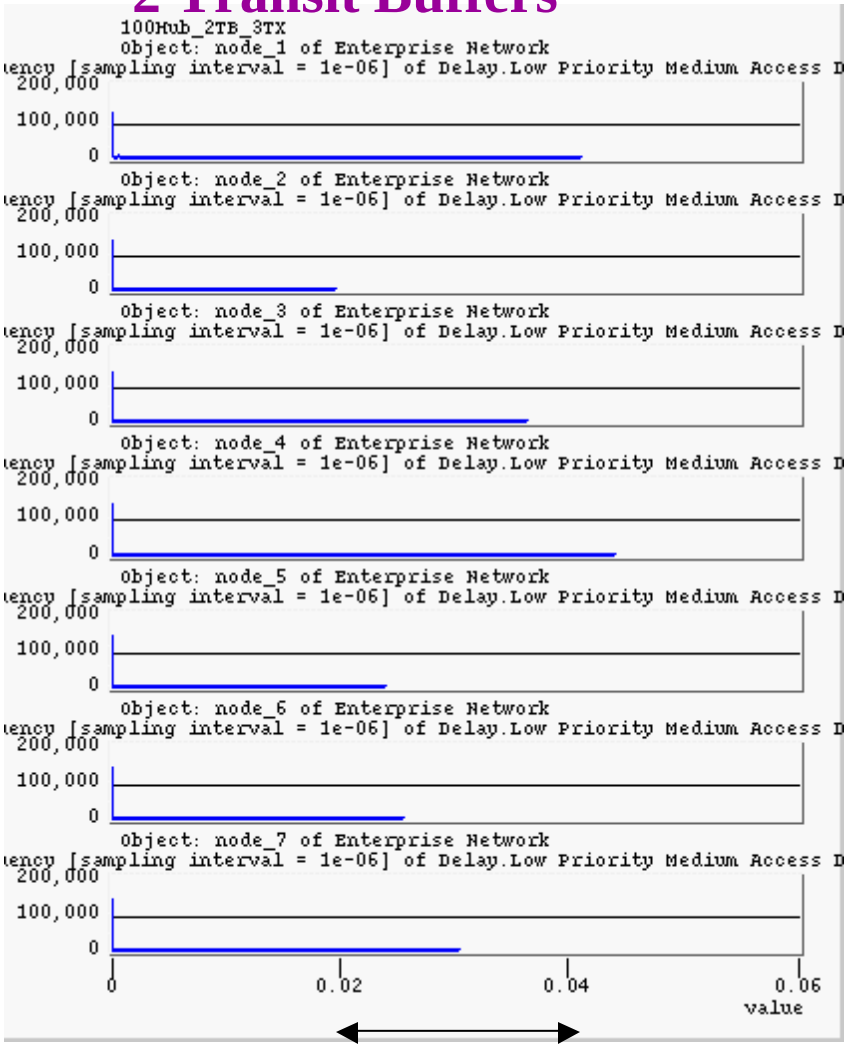

20msec

Conclusion

AuroraNetics, Inc.

- With 3 transmit buffers and a rate limiter for MP, we can provide guaranteed bandwidth to MP
- 2 transit buffers perform just as well in terms of throughput, delay and jitter for provisioned traffic as 3 transit buffers
- Delay and jitter guarantees are compromised with single transit buffer for provisioned traffic
  (http://grouper.ieee.org/groups/802/17/documents/presentations/may2001/nu_ctvst_02.pdf)

AuroraNetics, Inc.

Thank you!

Q & A