Ultra Low Queuing Delay for ALL Applications

http://riteproject.eu/dctth

IETF Journal: http://www.internetsociety.org/publications/ietf-journal-november-2015/ultra-low-delay-for-all

Koen De Schepper, Inton Tsang   Bell Labs

Olga Bondarenko, Bob Briscoe   [simula . research laboratory]

koen.de_schepper@alcatel-lucent.com

November, 2015

RITE
Reducing Internet Transport Latency

Super-Fast Internet?

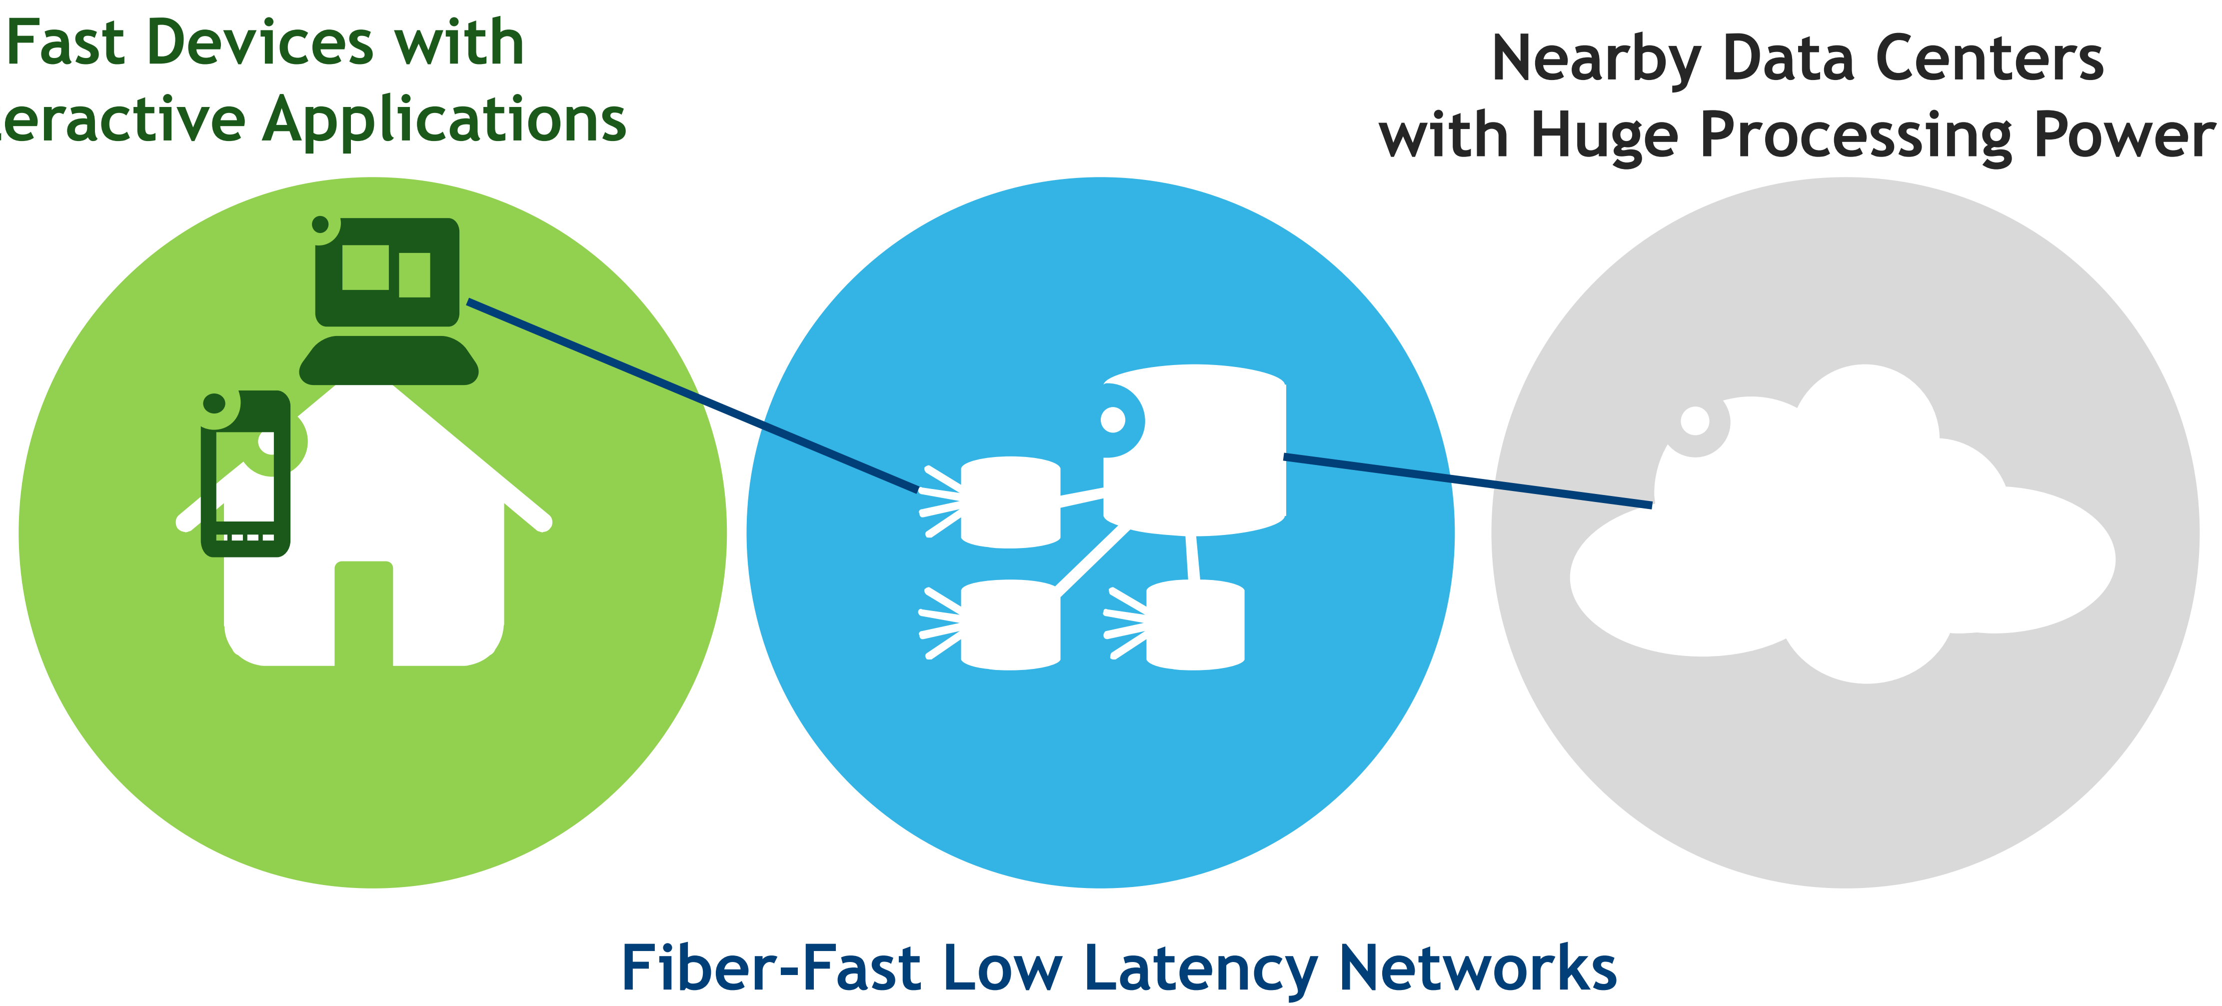

Super-Fast User Experience ???

Fast Devices with Interactive Applications

Nearby Data Centers with Huge Processing Power

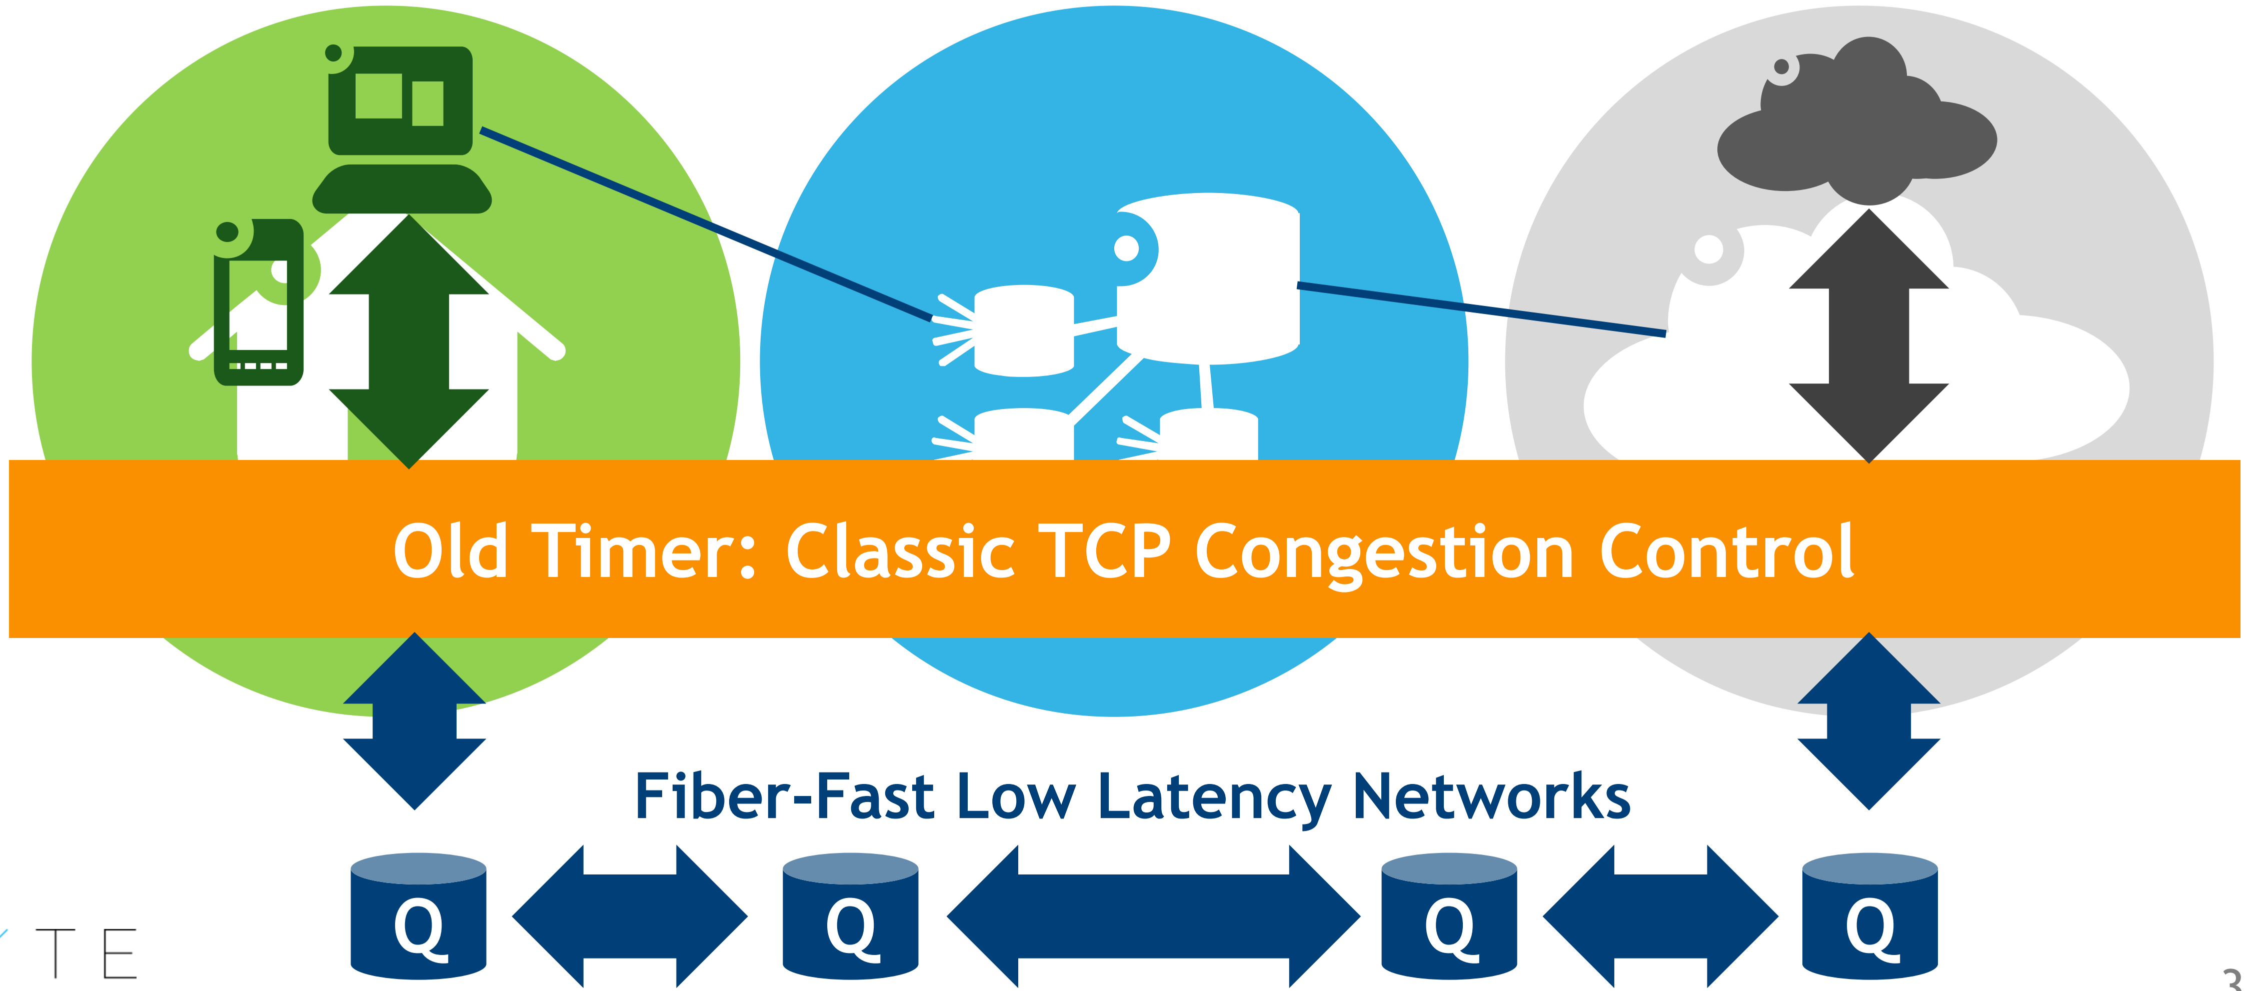

R/TE
Reducing Internet Transport Latency

Better TCP Exists, but is not compatible with the current Internet

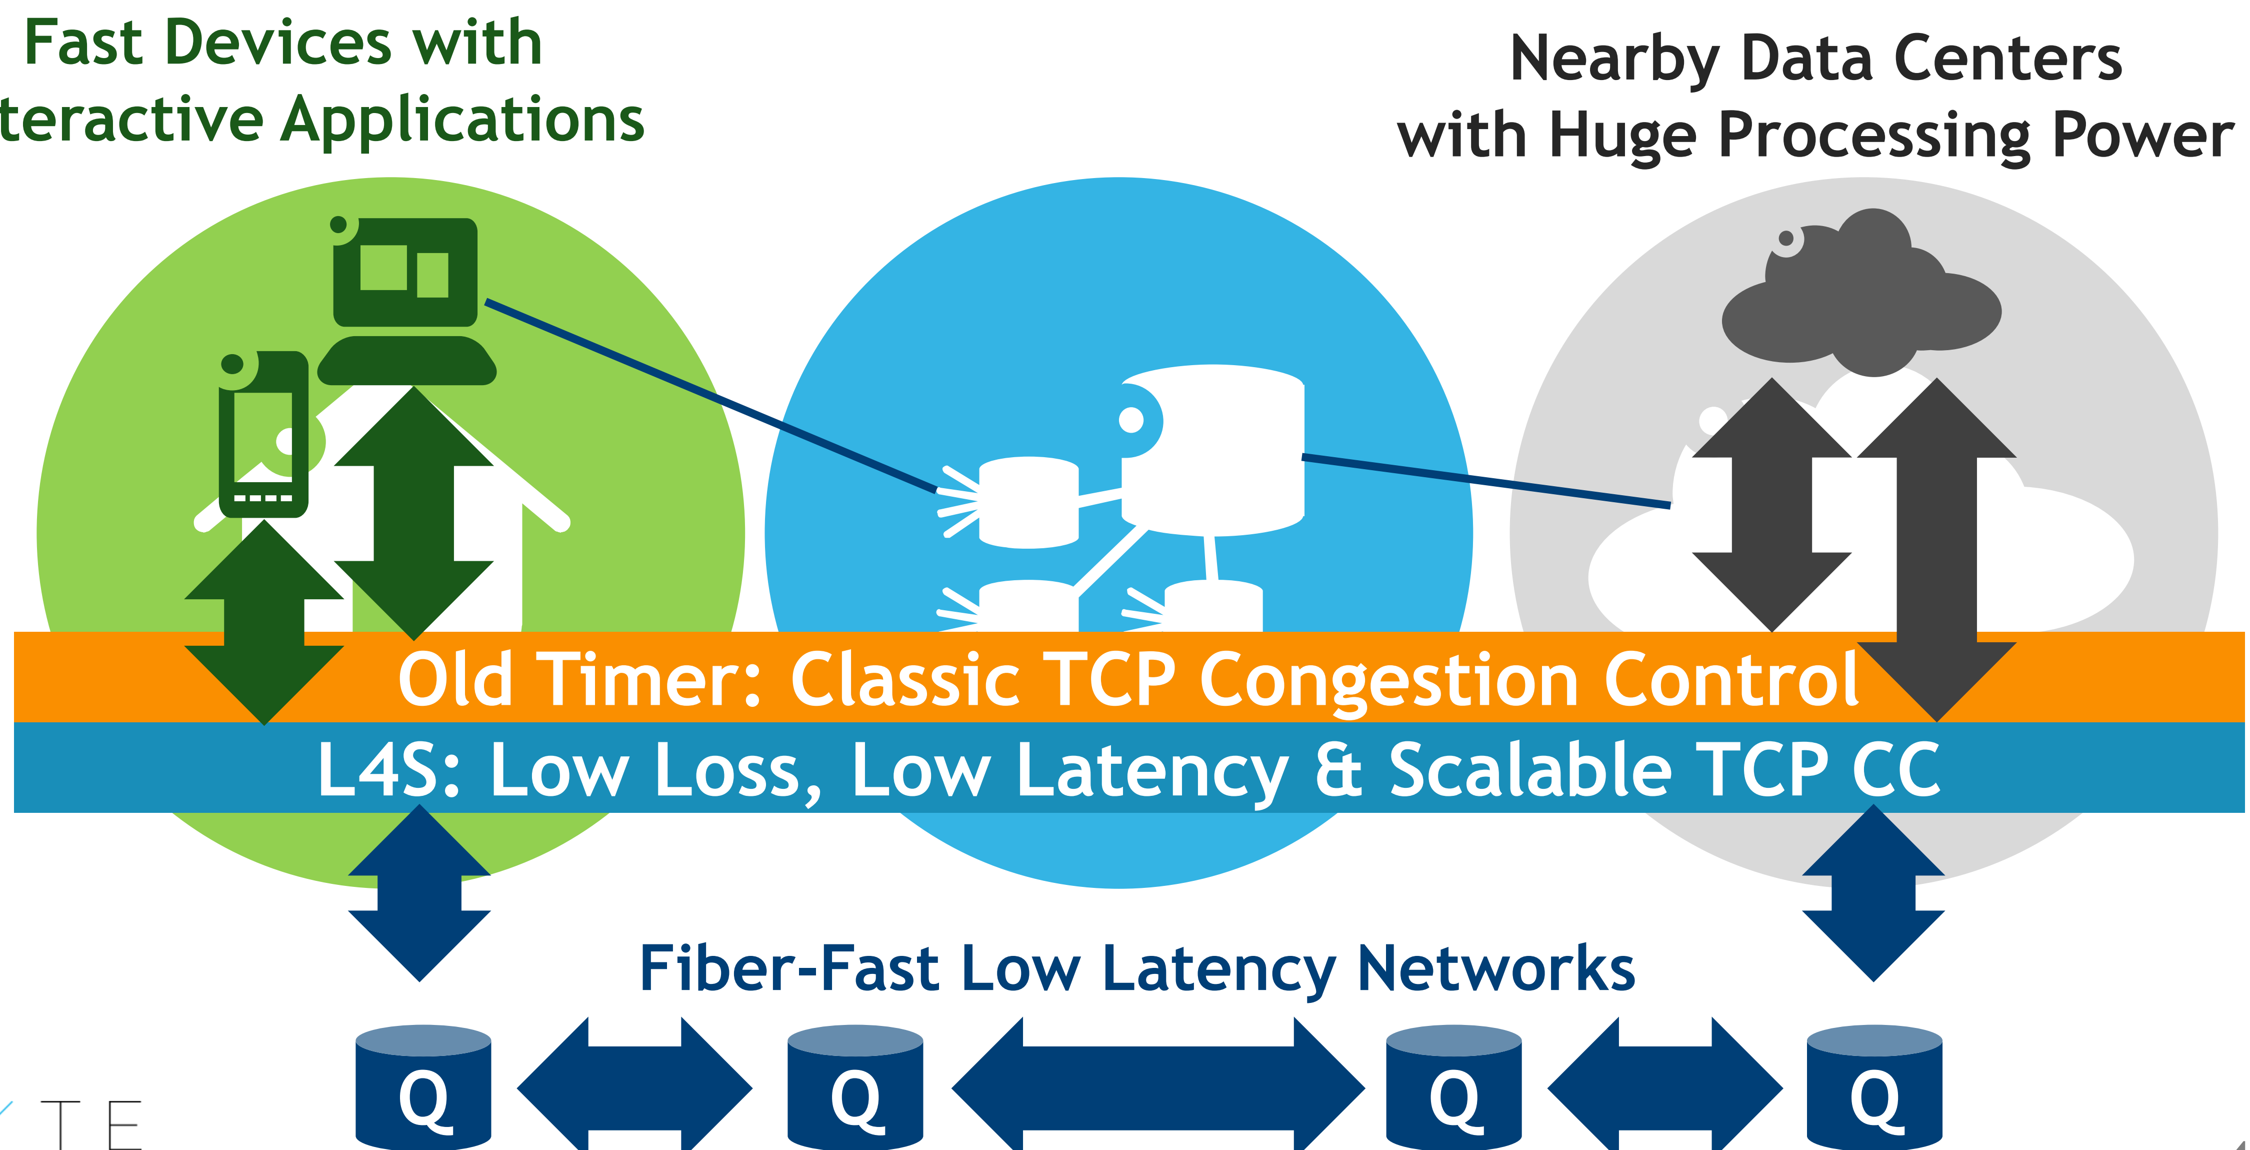

RITE
Reducing Internet Transport Latency

Comparison of Classic TCP and Scalable TCP Cloud-hosted Interactive Panoramic Video

We have high bandwidth

- but not predictable high speed & low latency

We will demonstrate the cause:

- The presence of ‘classic’ TCP (Reno and Cubic)
- The solution is ‘scalable’ TCP

Video link: <tbd>

Comparison of Classic TCP and Scalable TCP Cloud-hosted Interactive Panoramic Video

Demo:

- real broadband network
- real base round trip (7ms)
- heavy background traffic (6 parallel file downloads)
- DCTCP is available on Linux and Windows Server, used 'as is'

Not Diffserv

- not QoS at the expense of other traffic
- every packet of every app: ultra-low delay
- only IP header identifier, no transport header, no DPI
- only 2 queues, no FQ (queue per flow)

Also if all TCP traffic is scalable TCP

- there are no quality problems

Comparison of Classic TCP and a Scalable TCP
Cloud-hosted Interactive Panoramic Video

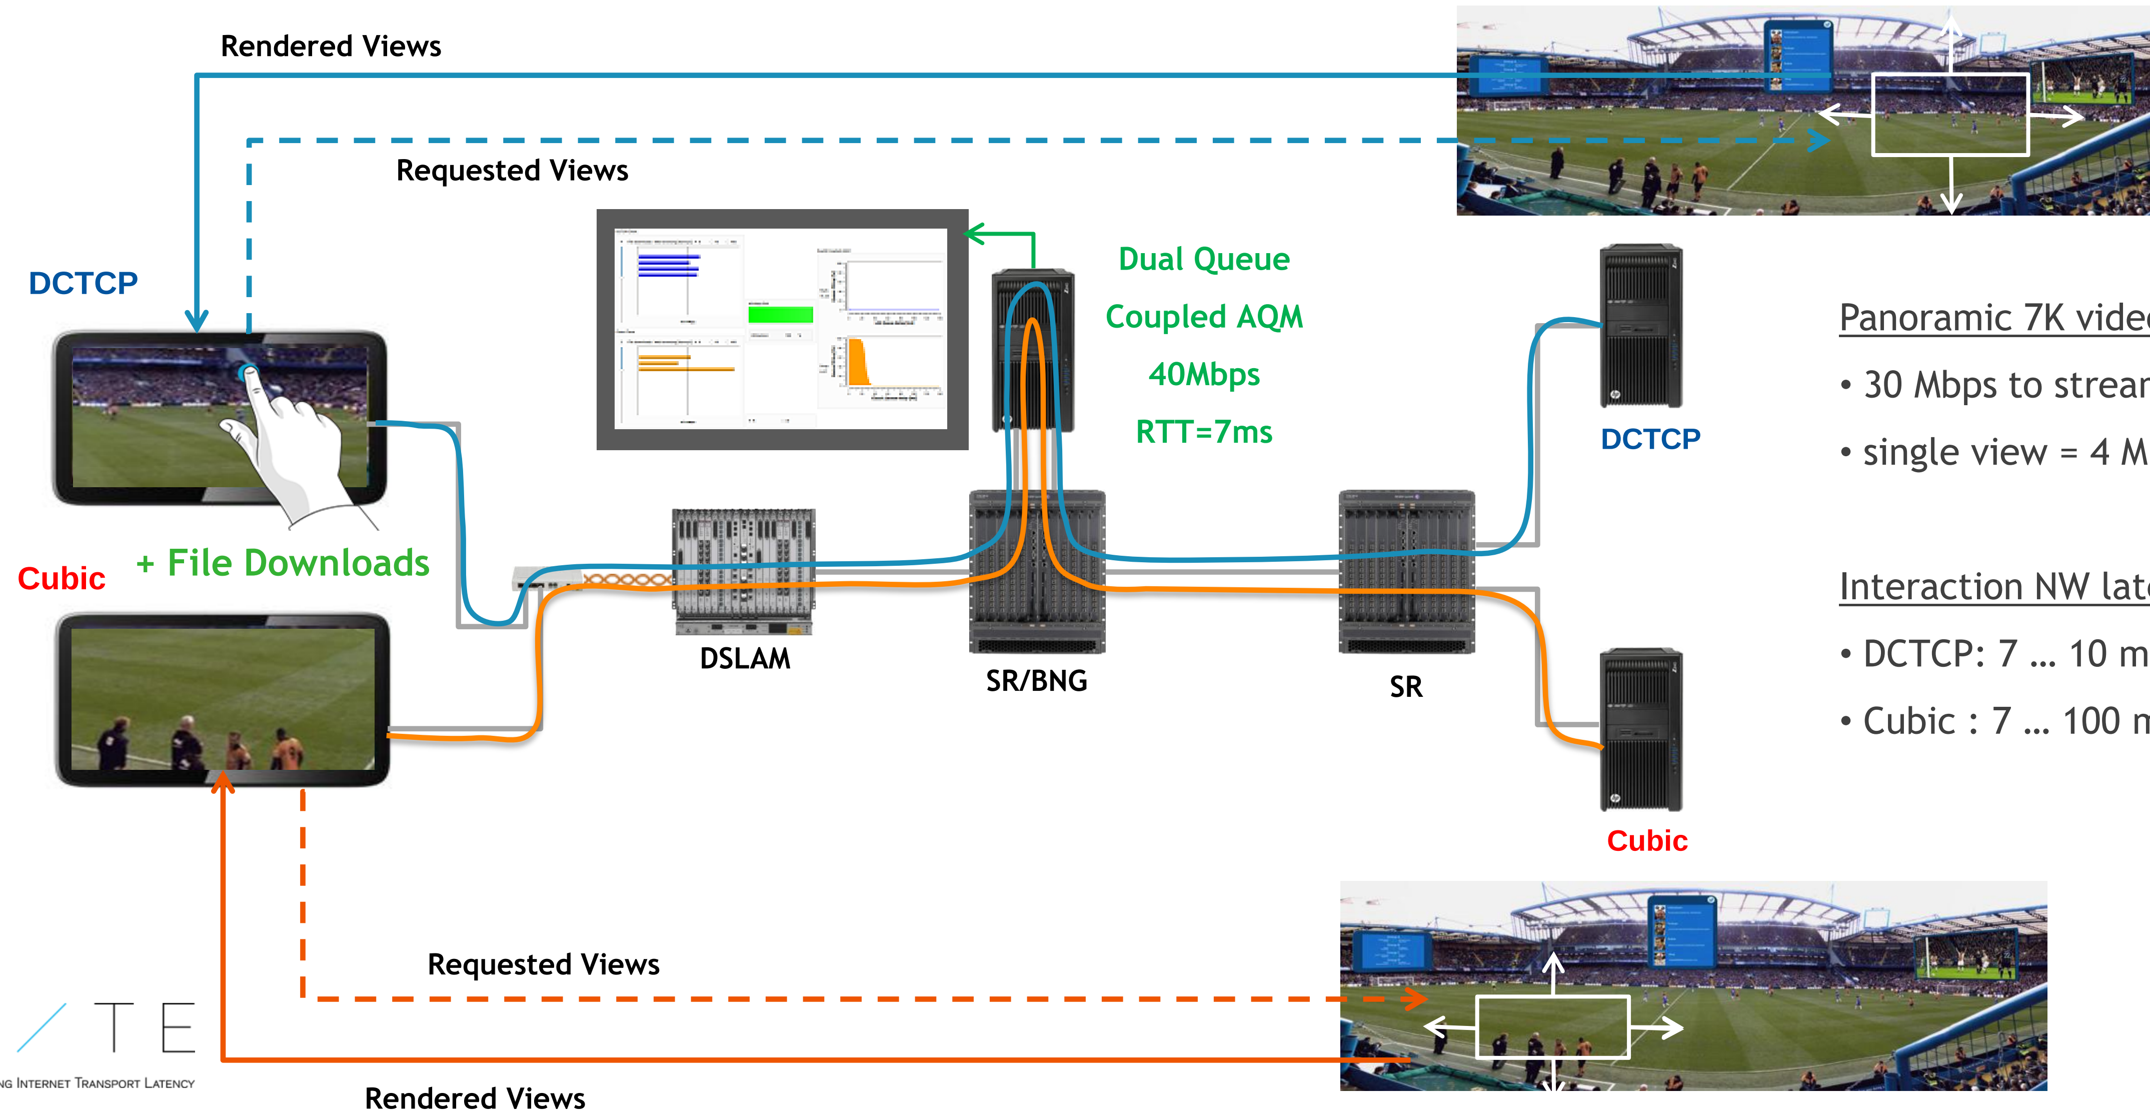

Panoramic 7K video

- 30 Mbps to stream
- single view = 4 Mbps

Interaction NW latencies:

- DCTCP: 7 … 10 msec
- Cubic : 7 … 100 msec

RITE
Reducing Internet Transport Latency

Why DISPATCH

We're not asking DISPATCH to dispatch something

We'll further show/explain the benefits to applications

But this involves

- 3 wgs in the transport area
- changes to hosts and network (bottlenecks) before it will bring benefits

Our question for DISPATCH, as 'customers' of the transport area:

- Are the benefits useful?
- Enough to overcome the chicken and egg deployment problem?

Simple Solution in the network:
DualQ AQM provides Low Latency Marking and Compatibility

Scalable-only AQM is very simple

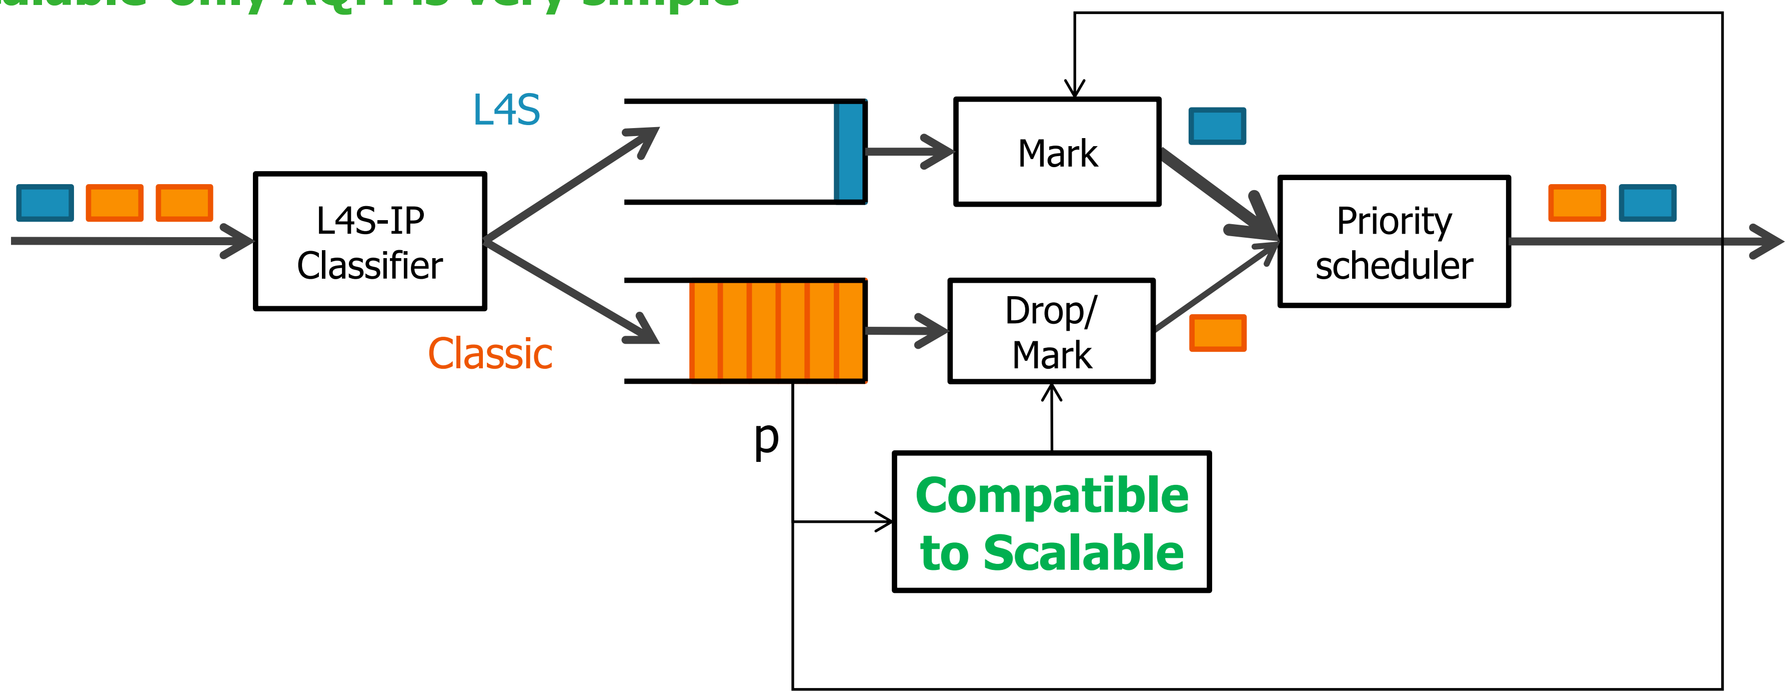

Coupled AQM needed for Compatibility

Allows Independent Migration

Demo on a Real BB Residential Testbed

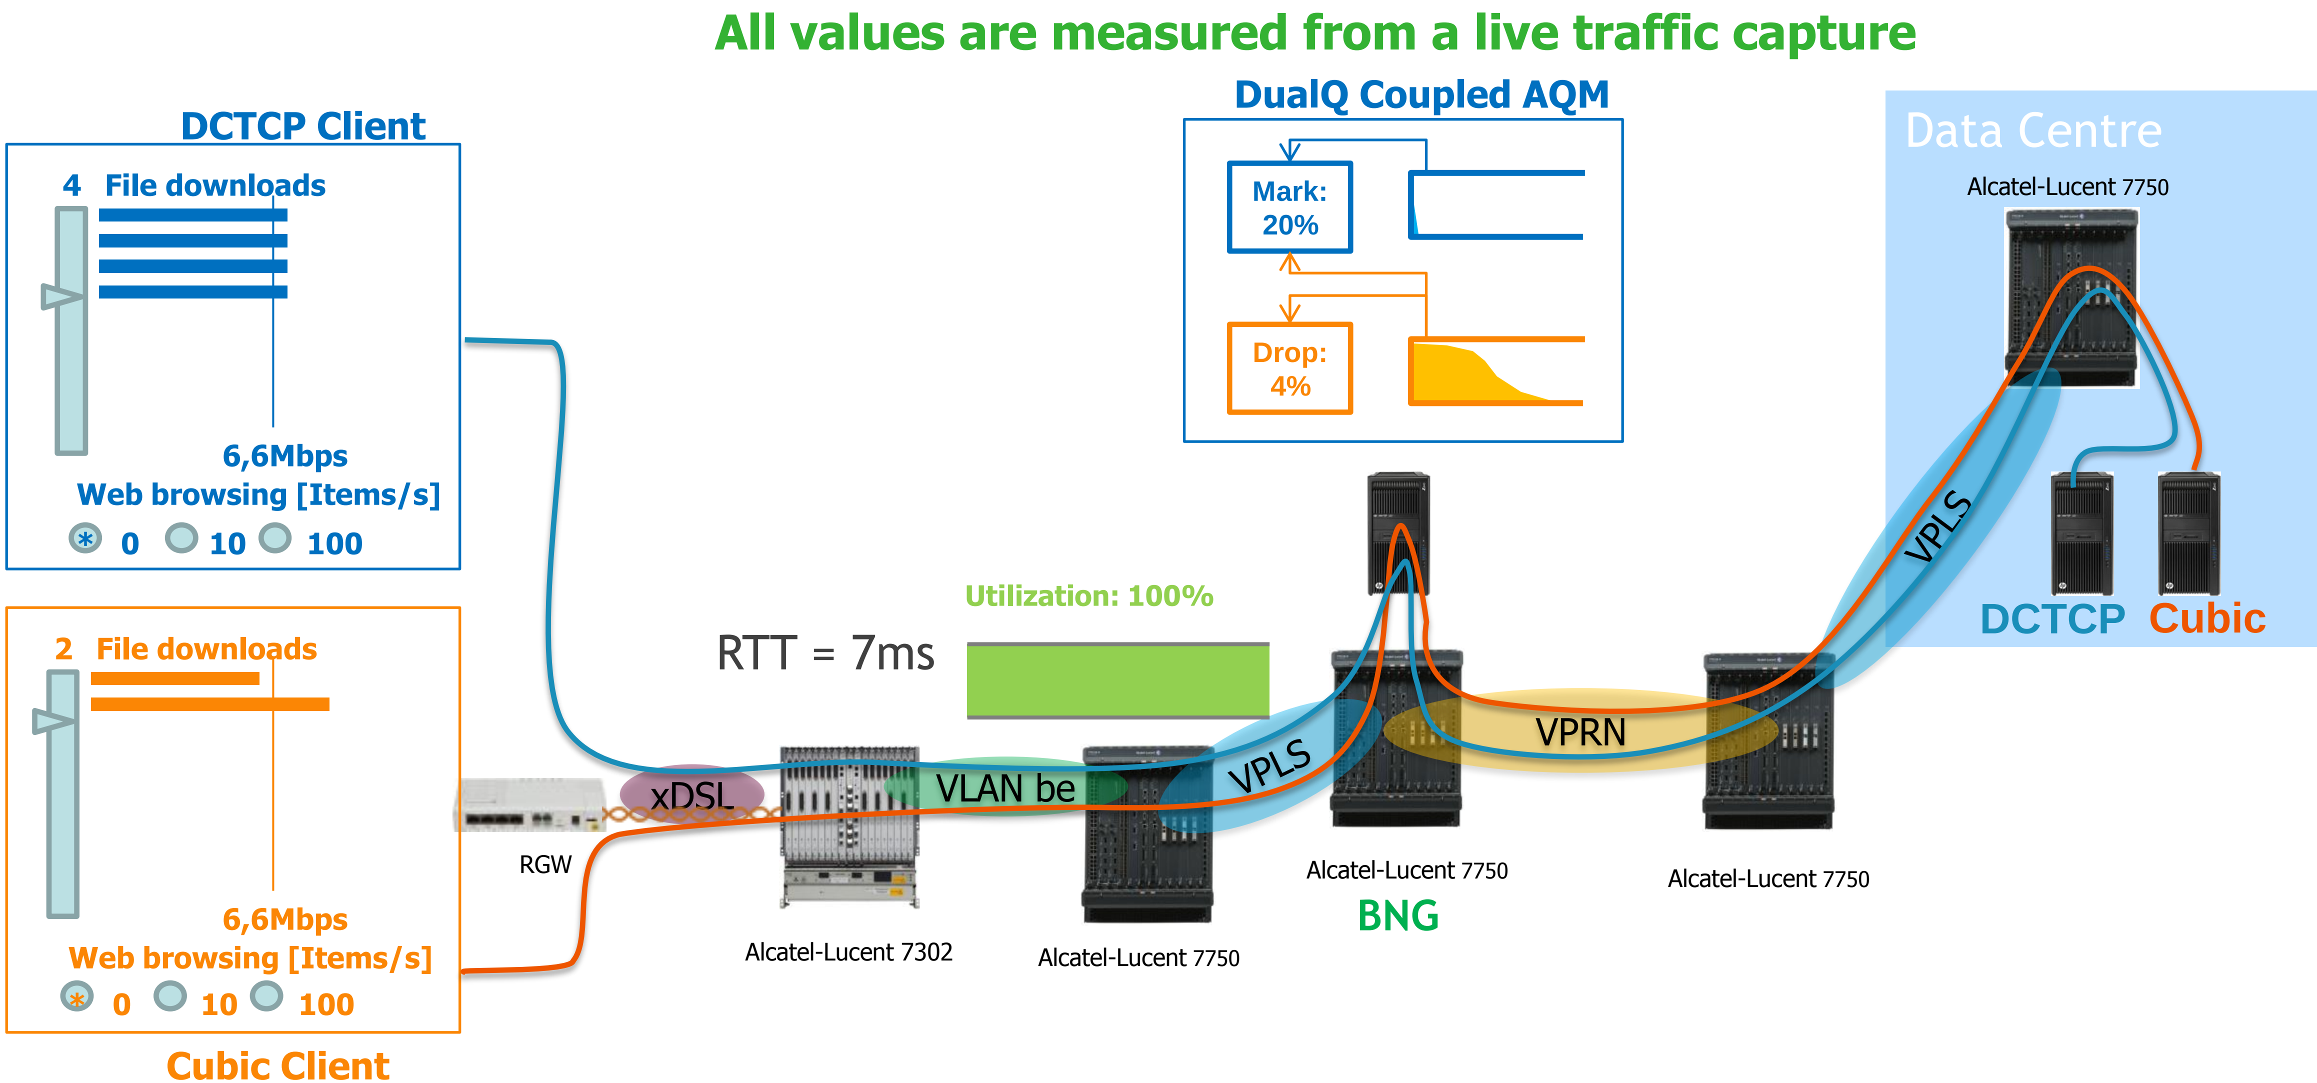

Demo on a Real BB Residential Testbed

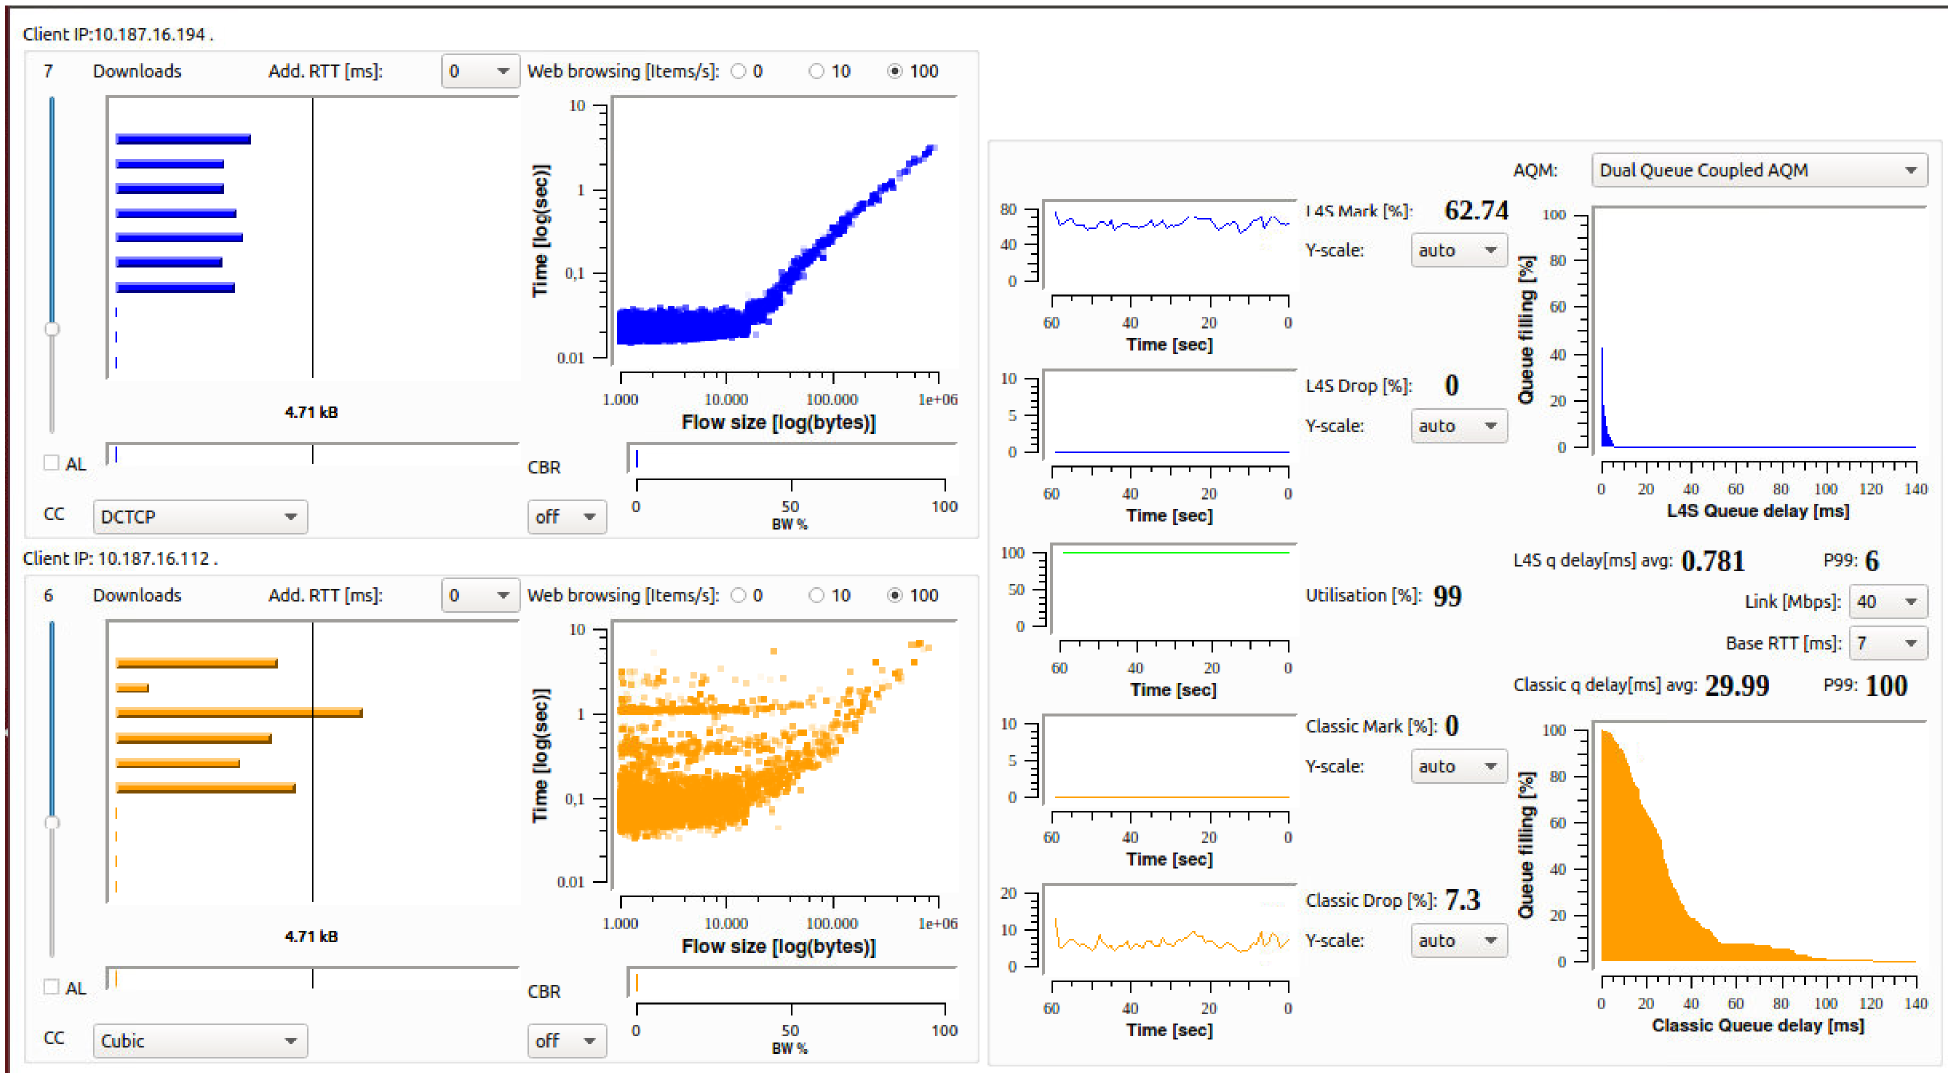

RITE
Reducing Internet Transport Latency

Improved Web browsing experience

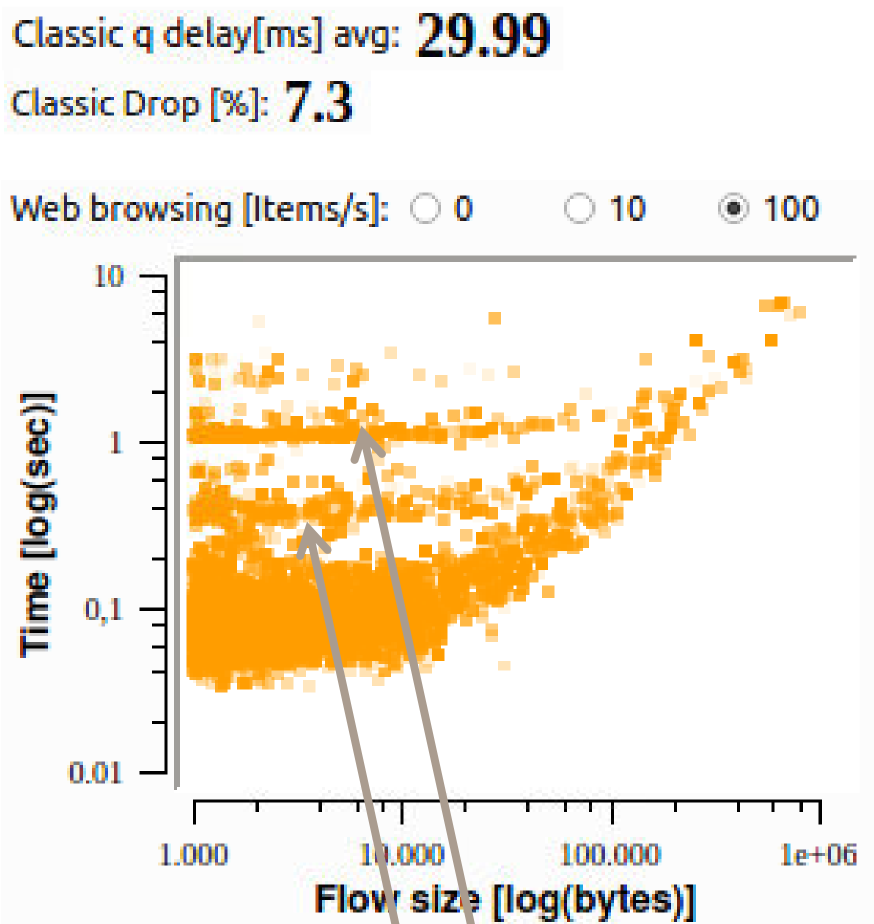

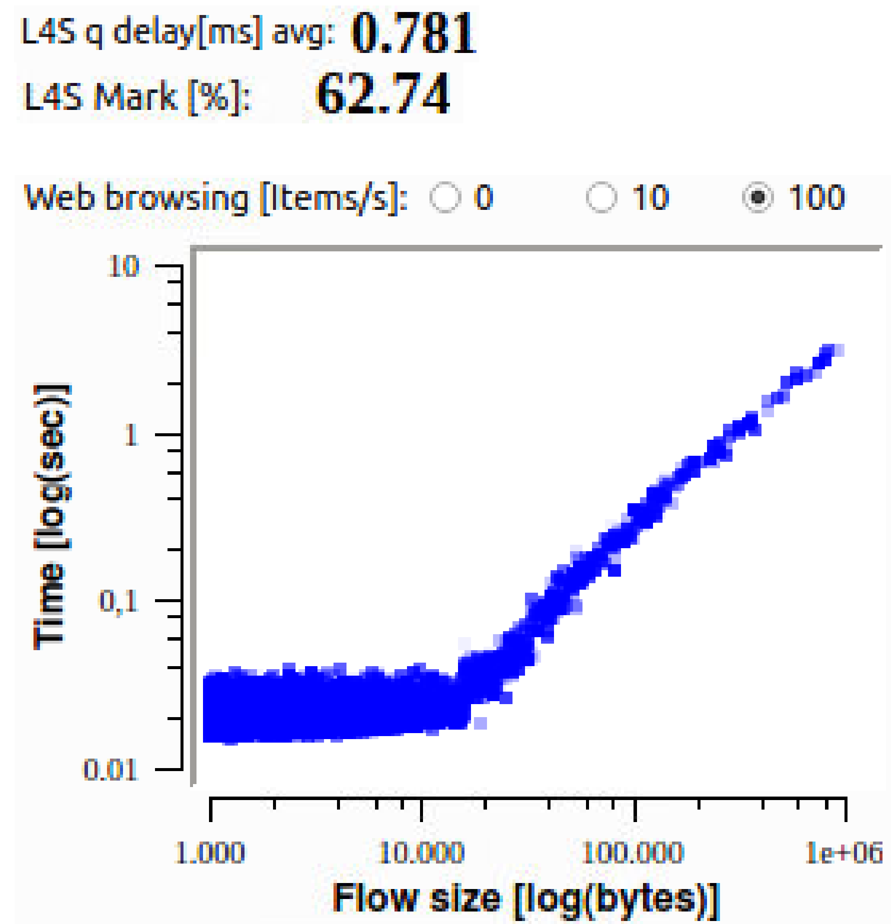

1 second timeout on lost connection setup packets

200ms timeout on lost last packets

HTTP Adaptive Streaming (HAS) Experiments

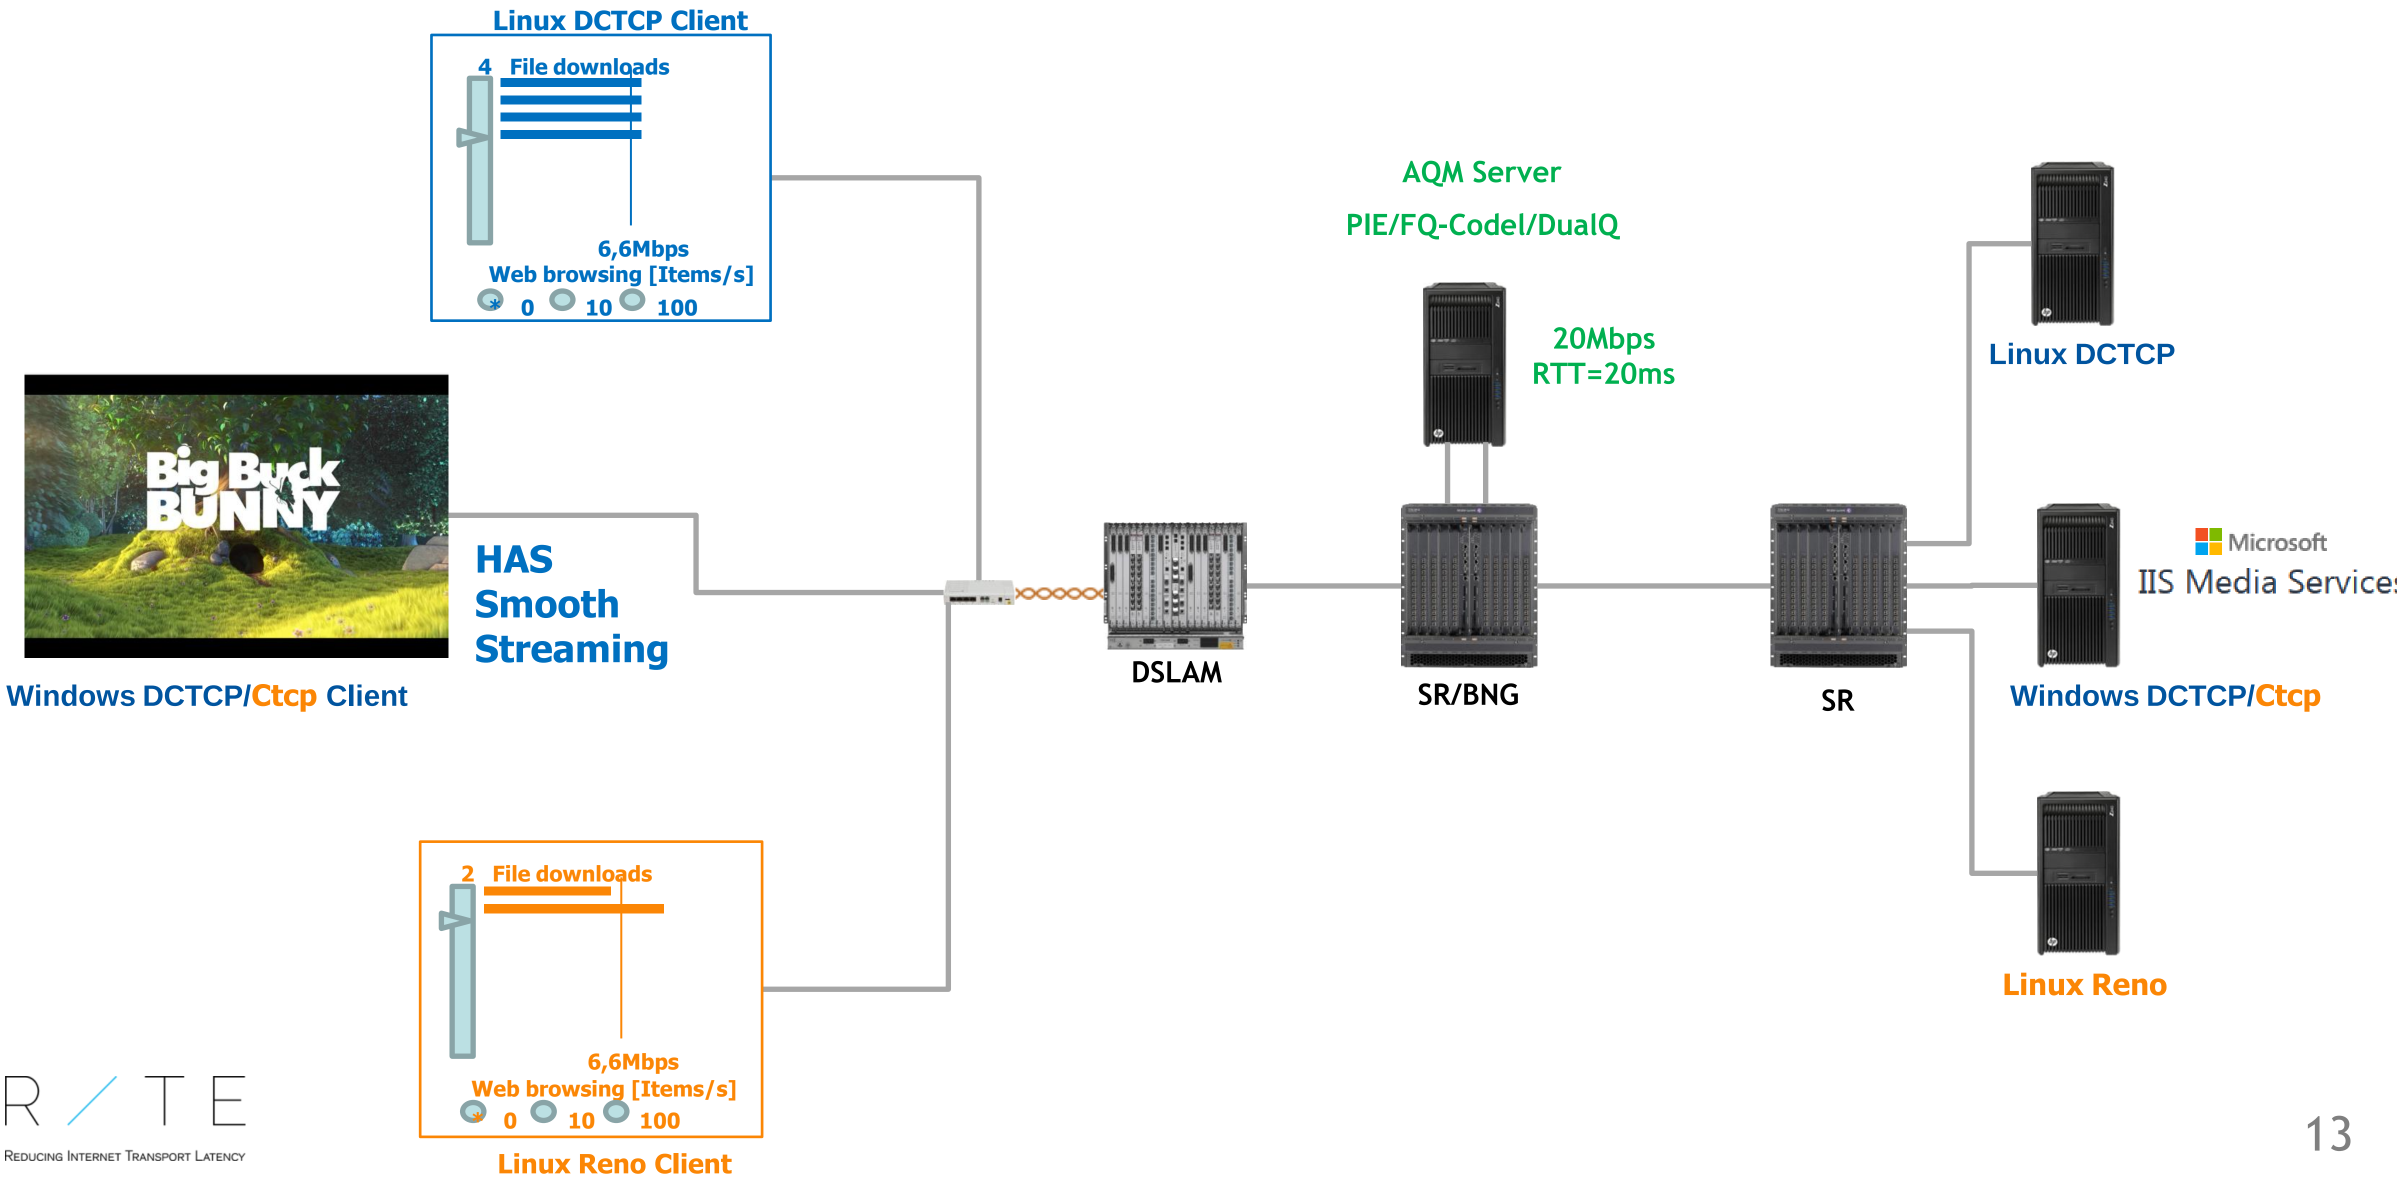

RITE
Reducing Internet Transport Latency

HTTP Adaptive Streaming (HAS) Experiments

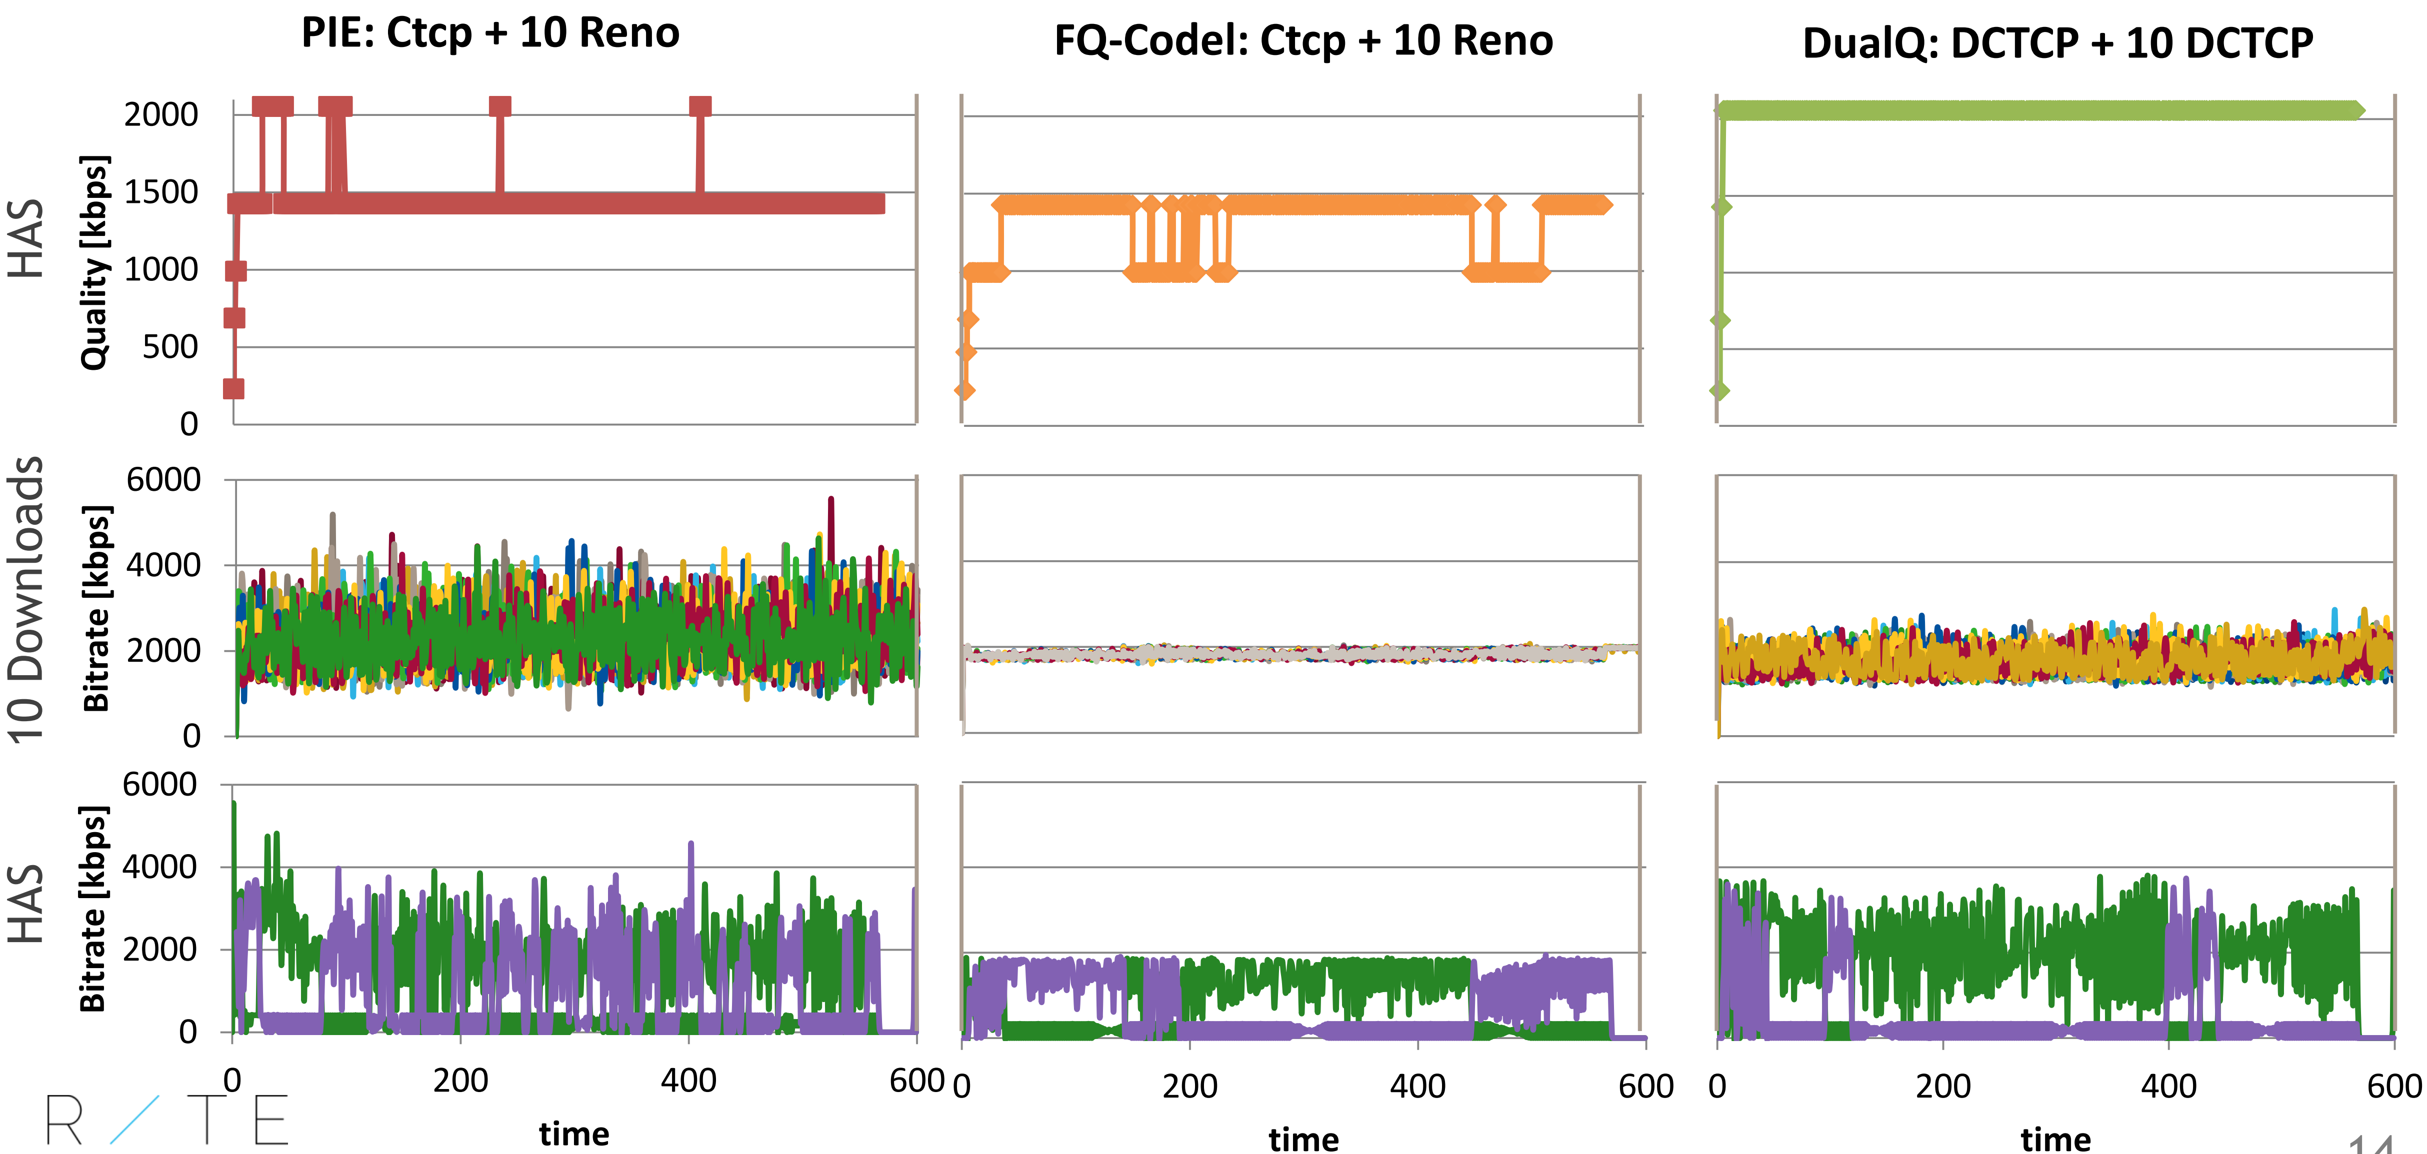

RITE
Reducing Internet Transport Latency

Using Just a Scalable TCP connection for Real-Time and Interactive Services?

Promising results

Prototyping / Evaluation by Applications

- DCTCP is available in Linux and Windows Server
- Collaboration
- Try it for your application…

koen.de_schepper@alcatel-lucent.com

ietf@bobbriscoe.net

Focus on Application
Not on Transport Functionality

R / T E
Reducing Internet Transport Latency

Questions

- Are the benefits useful?
- Enough to overcome the chicken and egg deployment problem?

- Support for adoption in the Transport Area?

Backup

DCTCP = Low Loss, Low Latency, and Throughput Scalability (L4S)

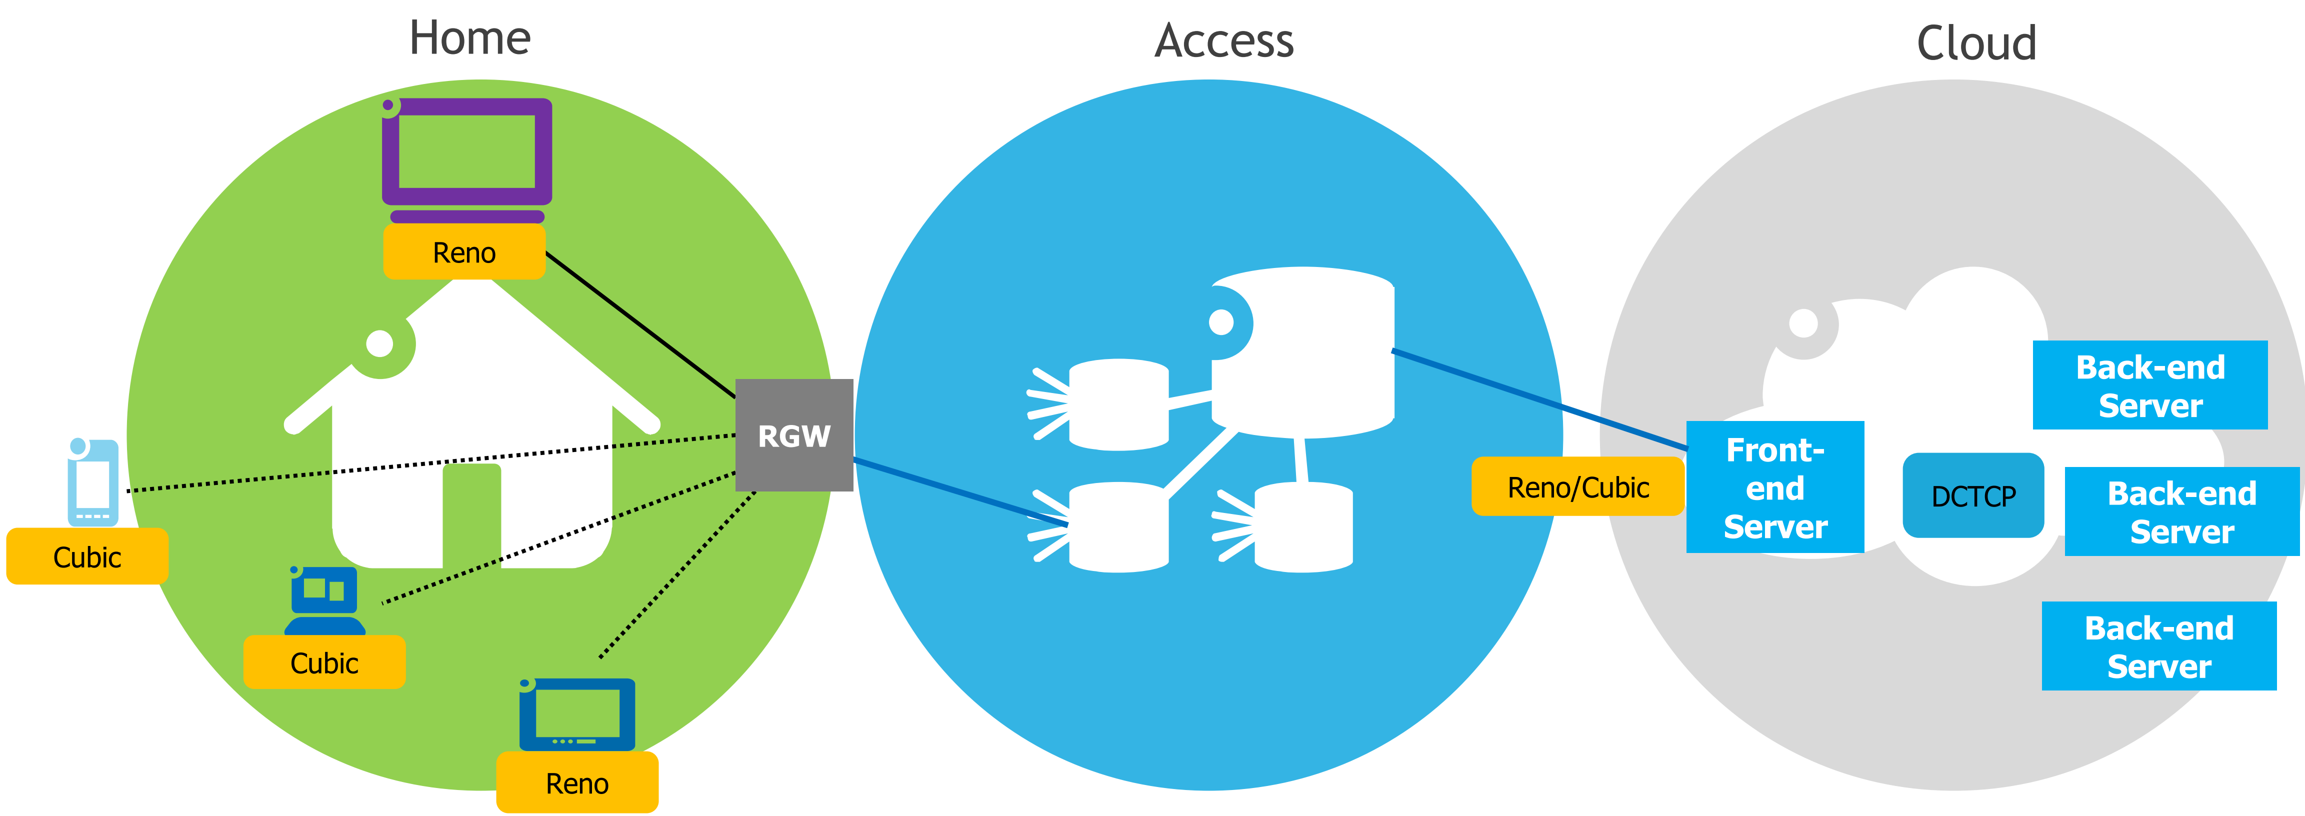

Large queues for high throughput and low drop
= Poor Latency
= Bad for interactive applications

ECN = No drop
ECN++ = Small queues
& Low latency & High throughput

R/TE
Reducing Internet Transport Latency

DCTCP to the HOME ?

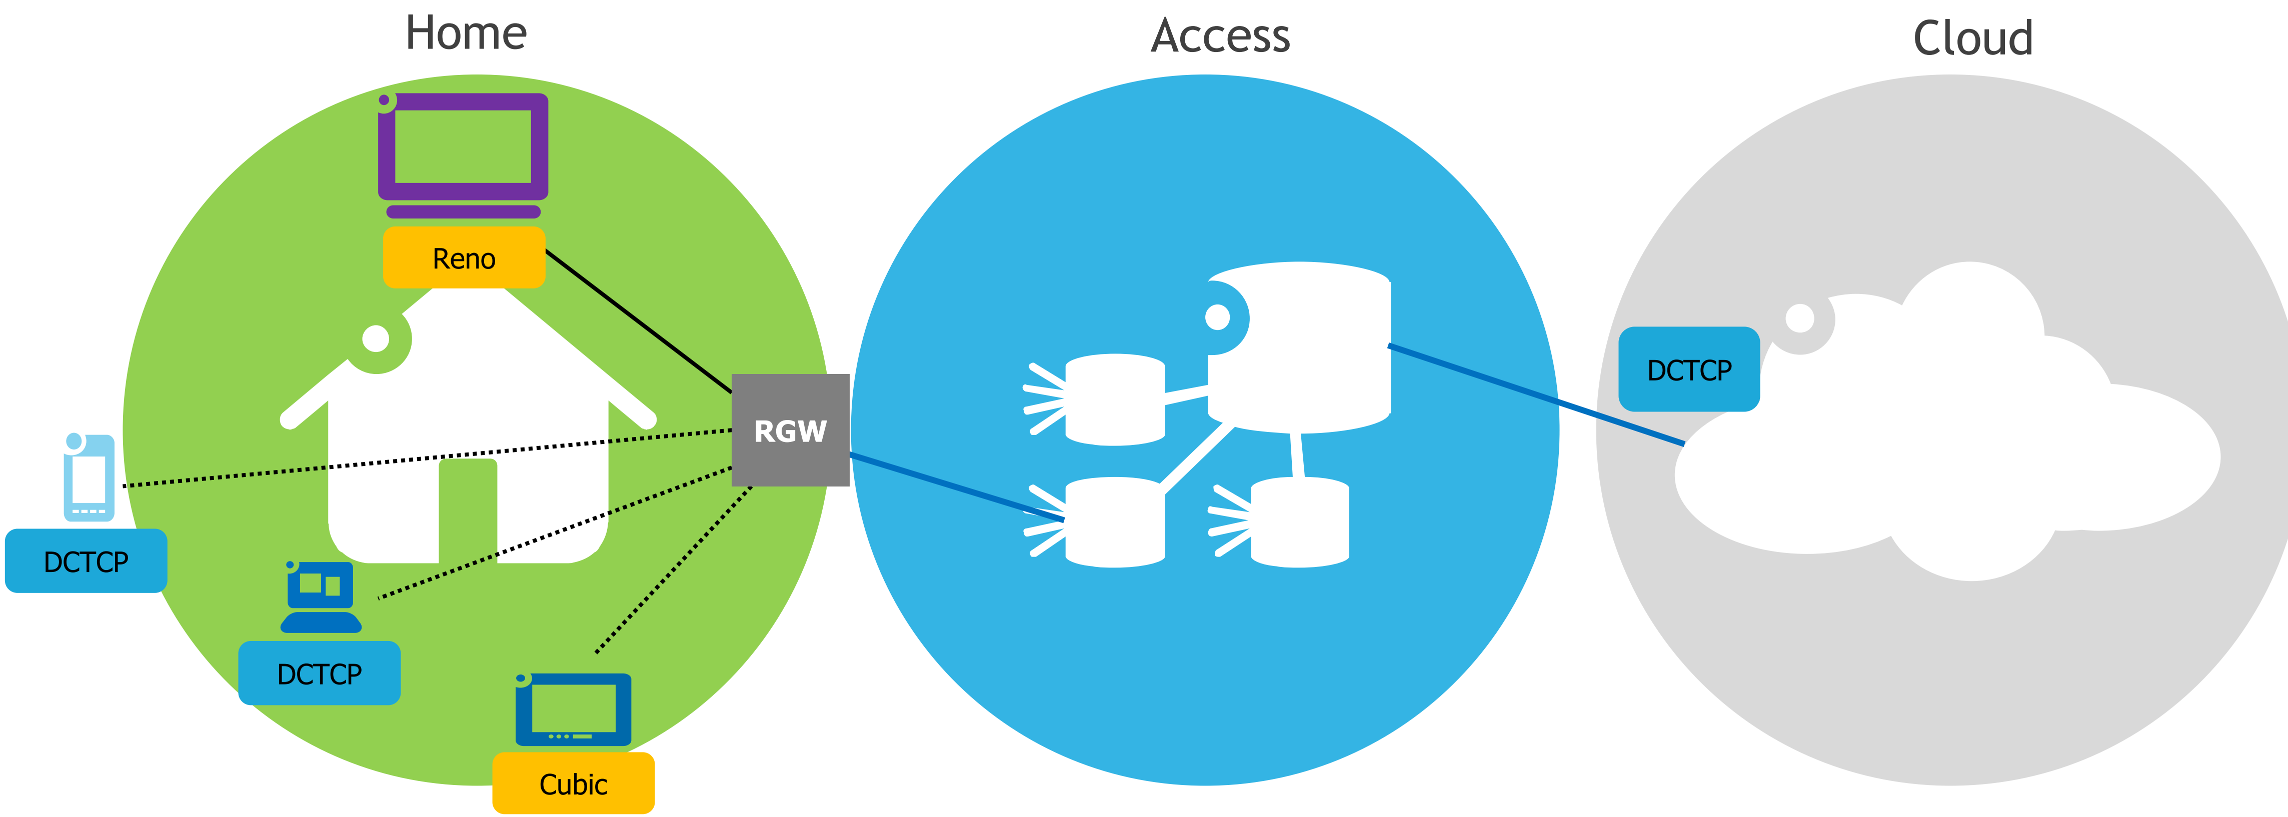

Windows and Linux 3.18
have DCTCP implementations ready

DCTCP available on
Windows Server and Linux 3.18
used internally in the data center

RITE
Reducing Internet Transport Latency

What can be done in the Network?

Small Queues
AQM

Low Utilization
Low RTT Fairness
High Drop
No Burst Resilience

Lower Latency ?

Very High Link Speeds

Lots of Memory in Qs
Very Low Drop
High Fairness Variations
Very Slow Up to Speed

Lower Latency ?

No Solution in the Network only

RITE
Reducing Internet Transport Latency

What can be done in the End-Systems?

Fix in TCP ? →

High Utilization
High RTT Fairness
Reliable Stable Throughput
No Drop
Burst Resilient
Fast Up to Speed
Very Lower Latency !
Not compatible with Classic TCP
Need for Network support

Fix in Applications ? →

Same constraints

No Solution in the End-Systems only

Dual Queue Coupled AQM
Concept 1: DualQ

DualQ to preserve low latency for L4S traffic

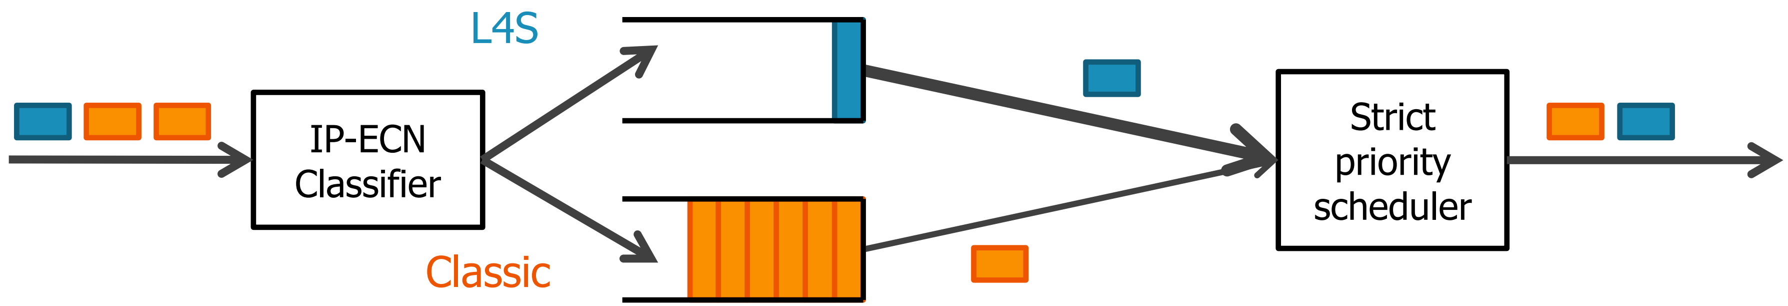

Dual Queue Coupled AQM
Concept 2: Coupled AQM

DualQ to preserve low latency for L4S traffic

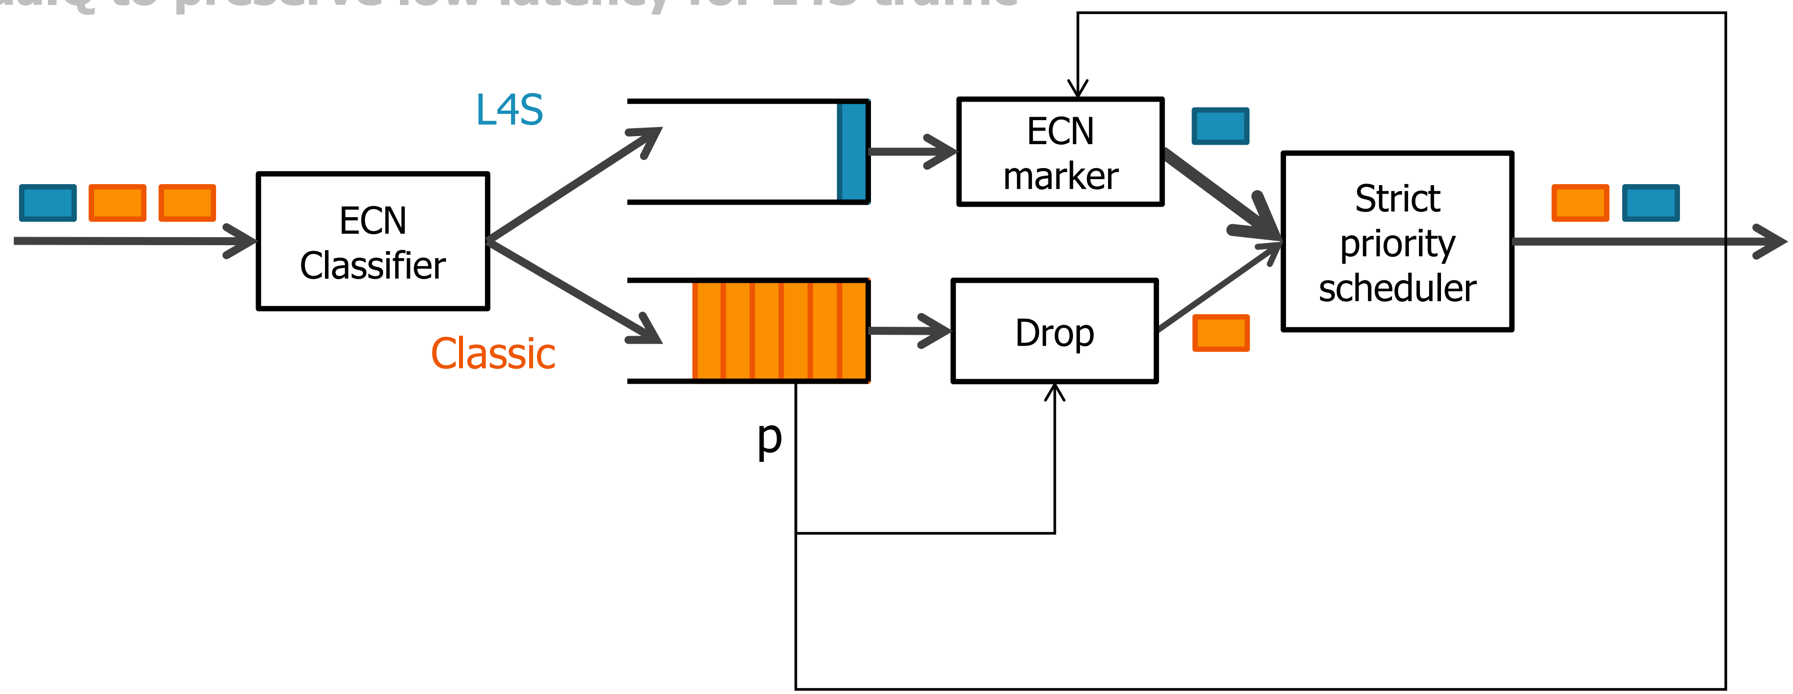

Coupled AQM to control priority traffic

Dual Queue Coupled AQM
Concept 3: Don’t Think Twice to mark

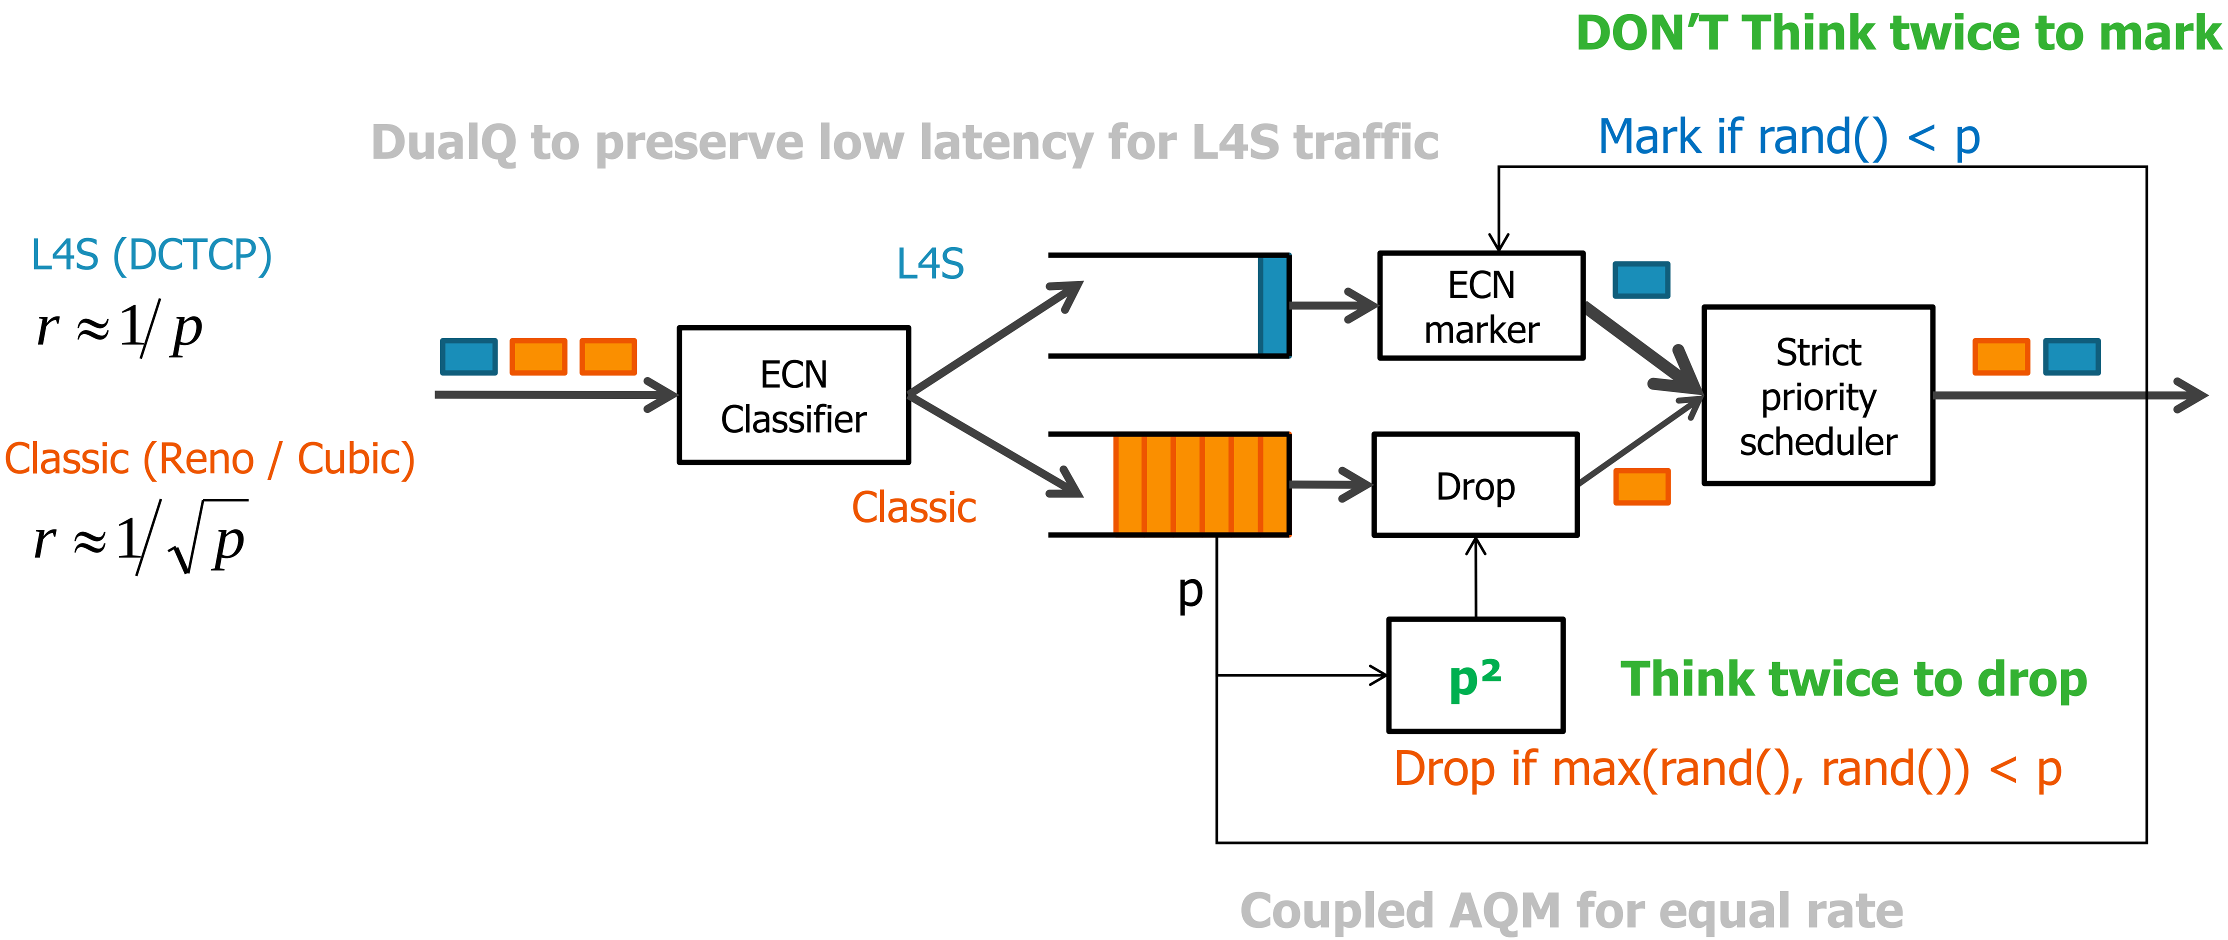

Detailed Implementation
2 Details

NO SMOOTHING for L4S

L4S AQM if no Classic traffic

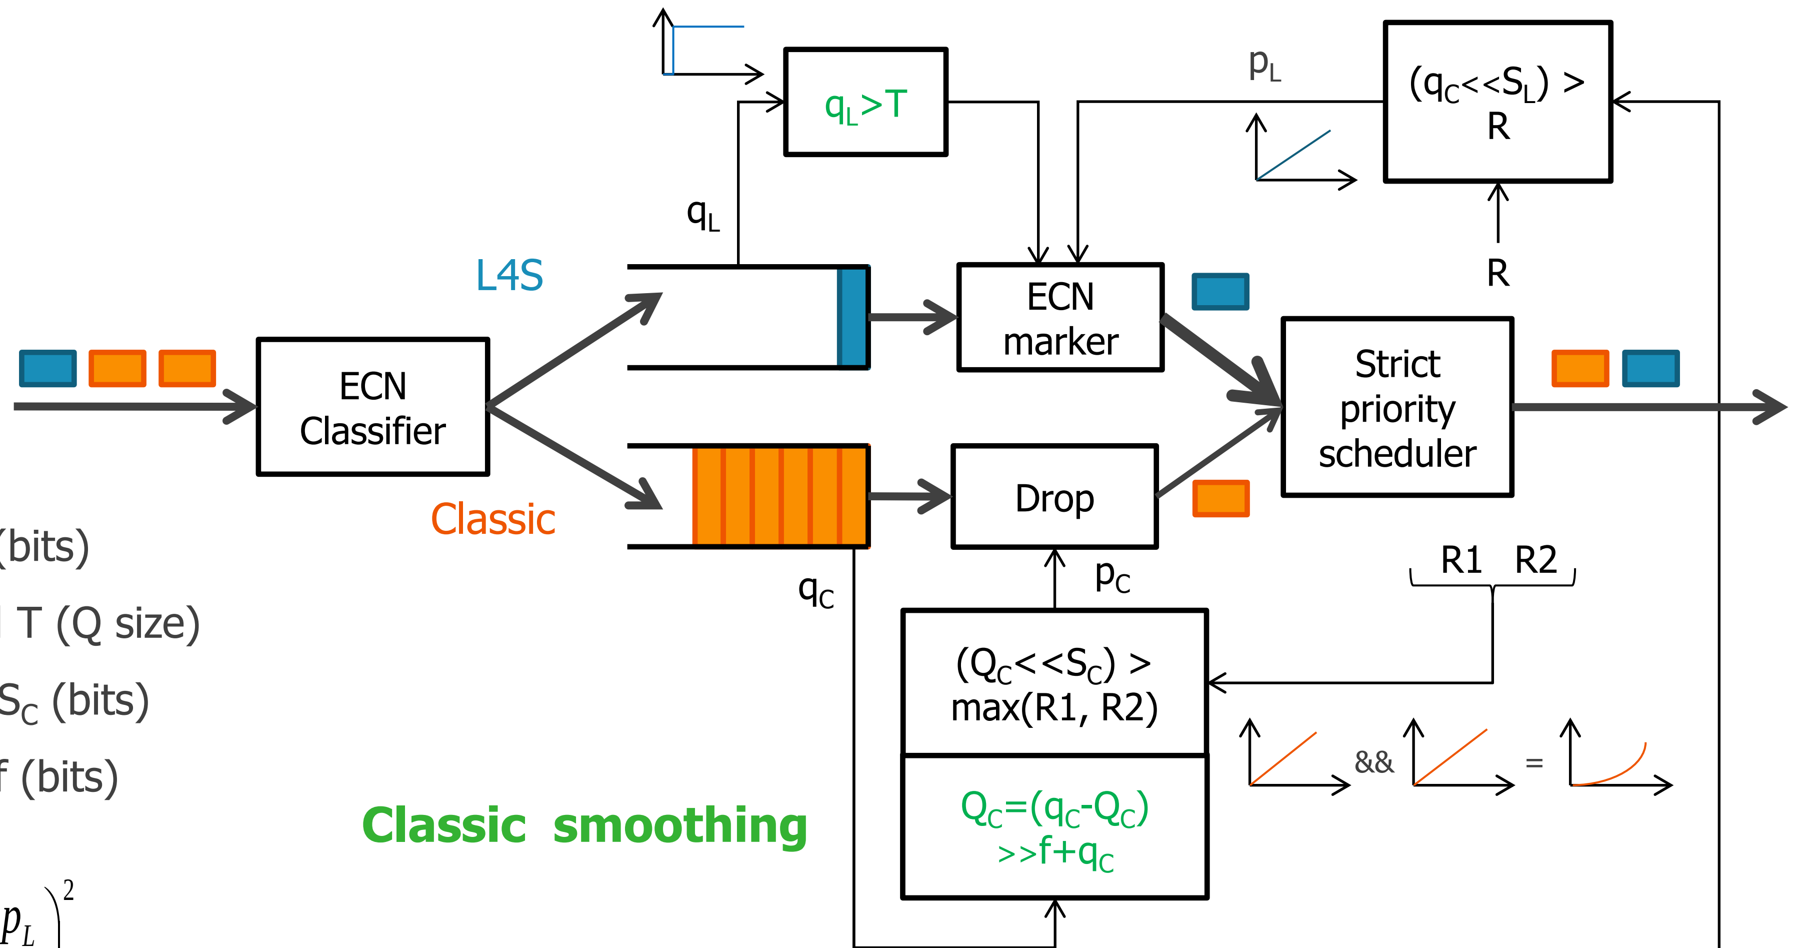

Classic smoothing

4 parameters:

- L4S slope S_L (bits)
- L4S threshold T (Q size)
- Classic slope S_C (bits)
- EWMA value f (bits)

$k = S_L - S_C$

Coupling: $p_C = \left(\frac{p_L}{2^k} \right)^2$

RITE
Reducing Internet Transport Latency

Demo on a Real BB Residential Testbed

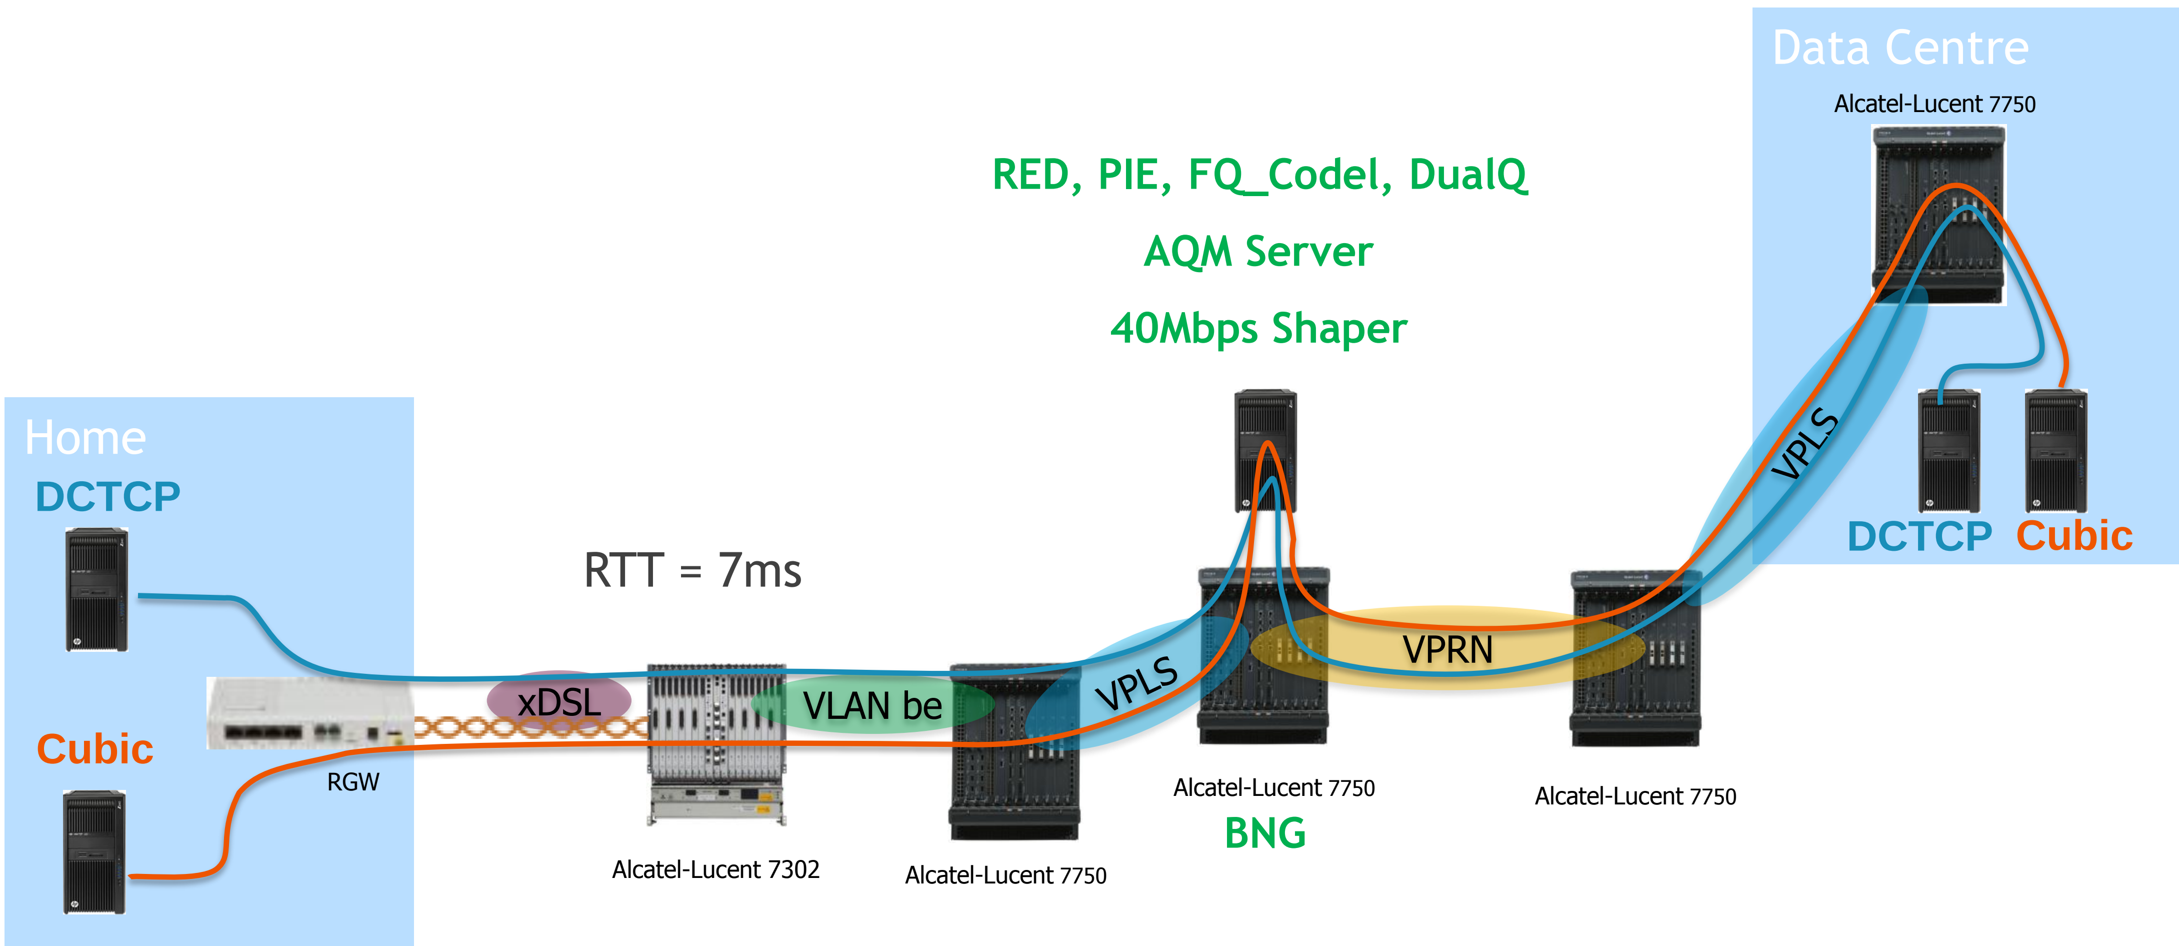

R / T E
Reducing Internet Transport Latency

Opportunity to support a new TCP CC family

Classic Congestion Controller Family

- Reno with $r \approx 1/\sqrt{p}$
- X RTTs per drop event
  - for Reno: X=45 RTTs on 40Mbps 20ms
    X= 5600 RTTs on 1Gbps 100ms
    → gets worse in the future (also for Cubic)
- Designed for Drop based networks in the 80'
- Was known not to scale to higher throughputs

L4S Congestion Controller Family

- L4S with $r \approx 1/p$
- C marked packets per RTT
  - (C=2 for DCTCP)
    → frequent feedback is better control!
- Design now for the future, using ECN better, to provide low loss, low latency and scalability
- DCTCP is big step forward and can be improved with incremental evolution

QUEUE SIZE AT DEQUEUE
1 TCP RENO FLOW (STEADY STATE)

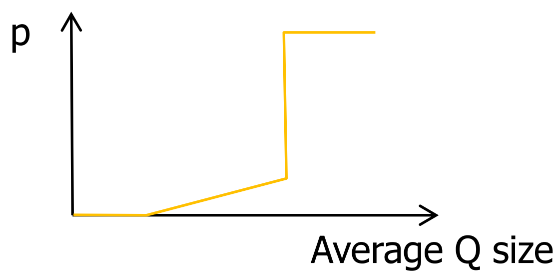

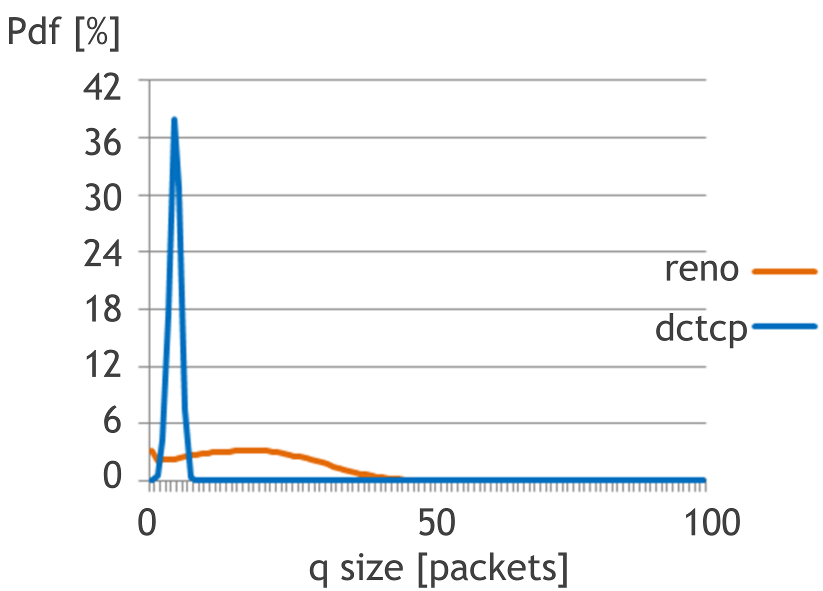

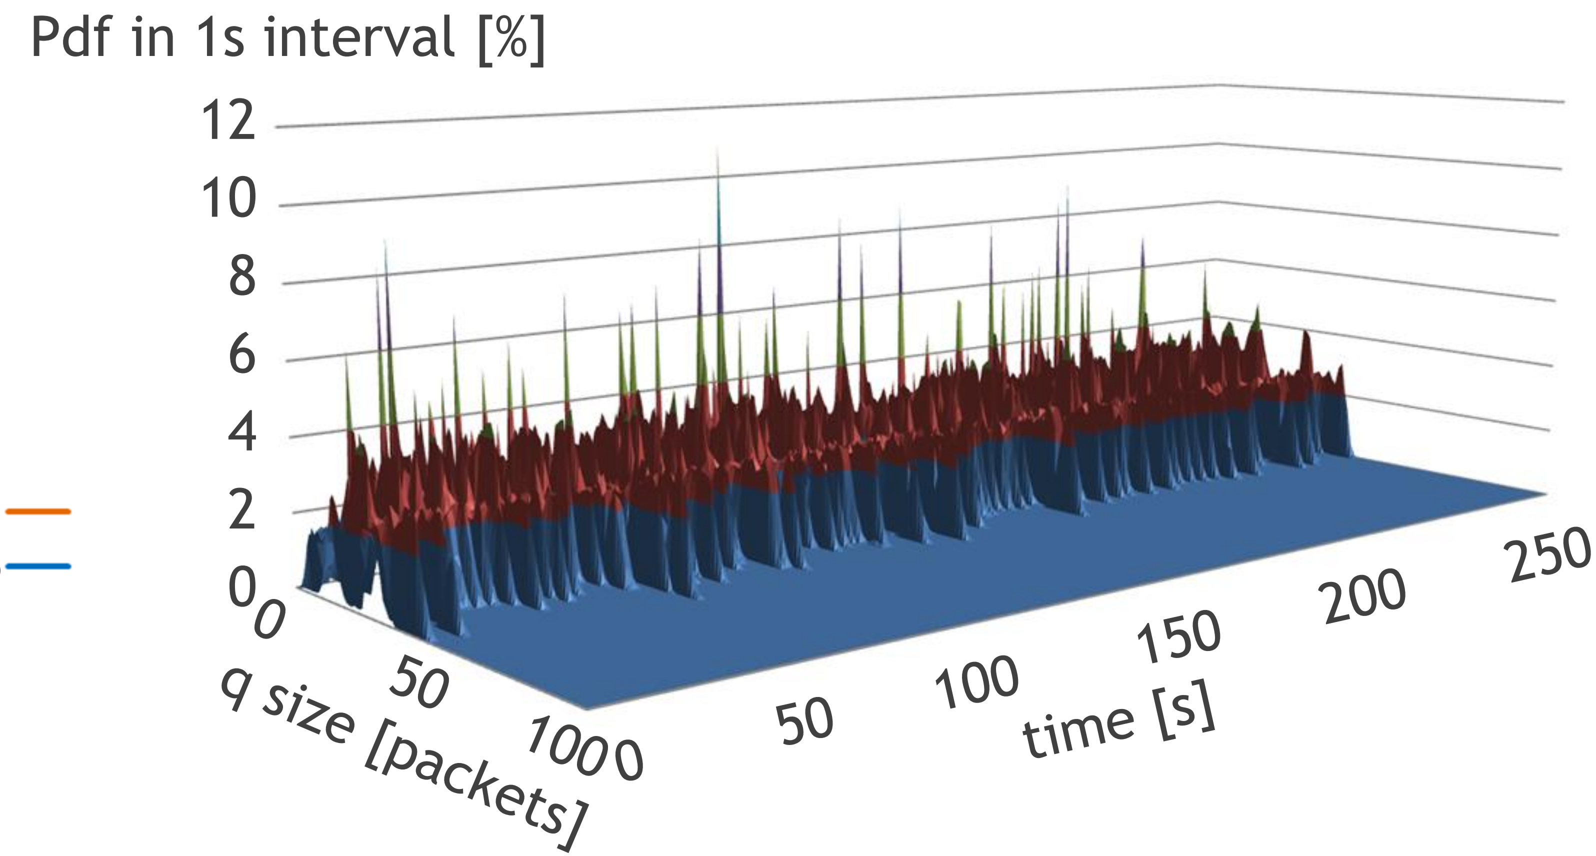

RITE
Reducing Internet Transport Latency

QUEUE SIZE AT DEQUEUE
1 DCTCP FLOW (STEADY STATE)

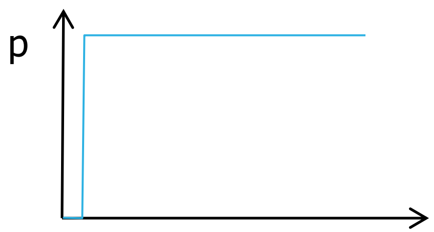

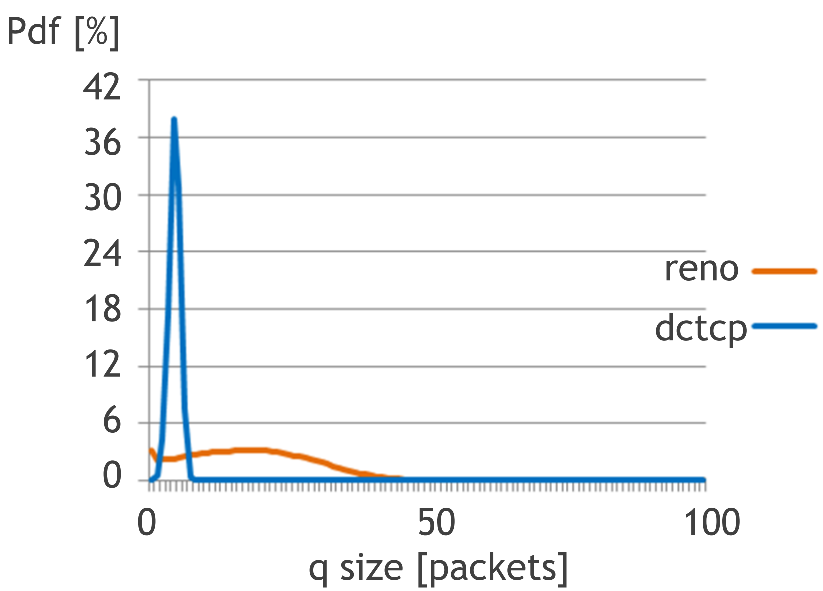

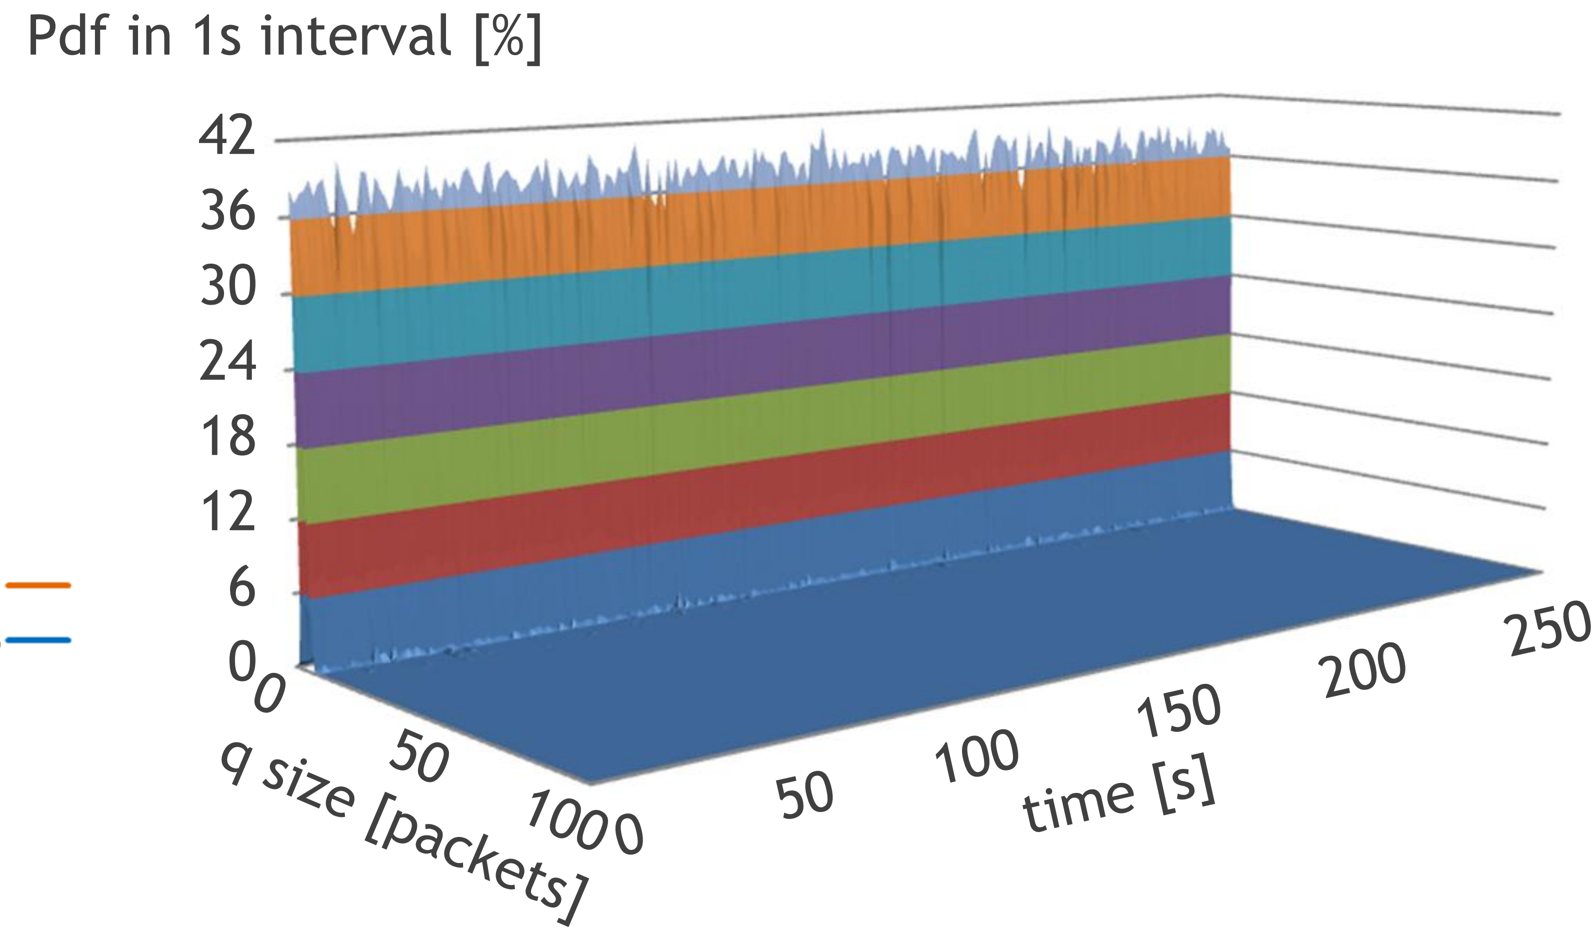

REDUCING INTERNET TRANSPORT LATENCY