

Fast Congestion management for Data Centers

draft-even-iccr-g-dc-fast-congestion-00

Roni Even
Rachel Huang



DC congestion control

- The use case that we are looking at is congestion control for Data Centers, a controlled environment, see RFC8085 section 3.6.
- Datacenter applications demand high throughput(40Gbps and above) with ultra-low latency of less than 10 microsecond per hop from the network, with low CPU overhead.
- Alternatives for network congestion direction can be classified as:
 - Based on estimation of network status: Traditional TCP, Timely, ...
 - Network provides limited information: DCQCN using only ECN, SCE,L4S, ...
 - Network provides some information: HPCC, ...
 - Network provides proactive control: RCP (Rate Control Protocol), ...

Proposed directions for DC congestion control

- Exploring these two directions
 - Reflect the network status more accurately – add metadata to the forward flow (e.g. using IOAM).
 - Notify the reaction point as soon as possible – report directly from the network to the sender (e.g. IOAM direct export)
- Issues to be addressed are discussed in the draft

Next Step

- Looking if there is interest in this direction.
- Side meeting Tuesday 8:30 Room VIP A
- Did some initial tests providing more information from the network. Will present initial test results in the side meeting.