



Contribution ID: 458

Type: **not specified**

Congestion Signaling (CSIG) for Linux TCP Data Center Networking

Saturday 13 December 2025 12:30 (30 minutes)

This proposal presents the overall design, challenges, and solutions for adding support for the Congestion Signaling (CSIG) protocol to the Linux networking stack. CSIG is being standardized within the UEC (<https://github.com/opencomputeproject/NET-UEC-CSIG>) and has broad support from prominent switch vendors.

CSIG is an in-band network telemetry protocol that allows end-hosts to obtain visibility into fine-grained network signals for congestion control, traffic management, and network debuggability. CSIG uses fixed-size L2 tags to collect metrics from network devices via a compare-and-replace mechanism, with data reflected from receiver to sender via transport layer options (e.g. TCP options). Additional reference: IETF draft (<https://www.ietf.org/archive/id/draft-ravi-ippm-csig-01.html>)

The talk will focus on topics such as:

1. Format and placement of CSIG tag and CSIG reflection header
2. Adapting Hardware and Software Generic Receive Offload (GRO) to handle CSIG-tagged packets.
3. Life of a CSIG-tagged packet in Linux TCP:
 - (a) End-host negotiation of CSIG feature during TCP connection establishment
 - (b) Receive path: L2 tag processing (reusing VLAN-handling logic), extraction of CSIG data, and CSIG data reflection back to the sender via TCP options on ACK packets
 - (c) Transmit path: CSIG L2 tag insertion with initial values.
 - (d) Use of the bottleneck link metrics in congestion control algorithms for making decisions on sending rate

Primary author: SHARMA, Sagarika

Co-authors: DUMAZET, Eric; CARDWELL, Neal; DE BRUIJN, Willem

Presenter: SHARMA, Sagarika

Session Classification: Networking Track

Track Classification: Networking Track