

Provide control of congestion

- for unicast traffic
- in networks with long-lived data flows with respect to bandwidth delay product within network domains of limited bandwidth delay product
- by means of control messages generated at the congestion point and received by the source of the data flow

Will not depend on any specific upper layer protocol.

Will be compatible with TCP/IP based protocols. There may be some TCP options that should not make use of congestion managed classes of service.

Support a bandwidth delay product of at least 1 Mbit, preferably 5 Mbit.

Allow for coexistence of multiple congestion-managed classes of service and of multiple non-congestion-managed classes of service. The data traffic to be controlled by the congestion notification mechanism will be segregated based upon the priority bits in the VLAN tag.

Support networks with full-duplex point-to-point links operating at a mix of link rates.

Define messages, congestion point behavior, reaction point behavior and managed objects.

Confine protocol messages to a domain composed solely of congestion notification capable bridges and end stations.

Consider inclusion of a discovery protocol (e.g. LLDP).

Will not introduce new bridge transmission selection algorithms or rate controls.

Per flow queuing and state will not be required in bridges.

The working group will coordinate with the Transport Area in the IETF on interactions with congestion-controlled Internet traffic, such as TCP, SCTP or DCCP.