

Requirements for Scalable Transport	Normative text	Apples opinion on the requirements No feedback on status/compliance	Evaluation/Remarks/Plans/Issues/Objections
1. Use of L4S Packet Identifier (A1.1)	<a href="#">Section 4.1</a> : A sender that wishes a packet to receive L4S treatment as it is forwarded, <b>MUST</b> set the ECN field in the IP header (v4 or v6) to the ECT(1) codepoint.		Currently we support setting the ECT(1) code point via Network.Framework API.
2. Accurate ECN Feedback (A1.2)	<a href="#">Section 4.2</a> : For a transport protocol to provide scalable congestion control it <b>MUST</b> provide feedback of the extent of CE marking on the forward path.		
(Scalable CC requirement)	<a href="#">Section 4.3</a> : As a condition for a host to send packets with the L4S identifier (ECT(1)), it <b>SHOULD</b> implement a congestion control behaviour that ensures that, in steady state, the average time from one ECN congestion signal to the next (the 'recovery time') does not increase as flow rate scales, all other factors being equal.		
(ECT(1) use needs Prague compliance)	<a href="#">Section 4.3</a> : In order to coexist safely with other Internet traffic, a scalable congestion control <b>MUST NOT</b> tag its packets with the ECT(1) codepoint unless it complies with the following bulleted requirements.		
(Prague compliance description)	<a href="#">Section 4.3</a> : The specification of a particular scalable congestion control <b>MUST</b> describe in detail how it satisfies each requirement, and for any non-mandatory requirements, it <b>MUST</b> justify why it does not comply.		
3. Fall back to Reno-friendly congestion control on packet loss (A1.3)	<a href="#">Section 4.3</a> : As well as responding to ECN markings, a scalable congestion control <b>MUST</b> react to packet loss in a way that will coexist safely with a TCP Reno congestion control [ <a href="#">RFC5681</a> ].		<b>MUST</b> comply seems right and this would be easy to implement.
4. Fall back to Reno-friendly congestion control on classic ECN bottlenecks (A1.4)	<a href="#">Section 4.3</a> : A scalable congestion control <b>MUST</b> implement monitoring in order to detect a likely non-L4S but ECN-capable AQM at the bottleneck. On detection of a likely ECN-capable bottleneck it <b>SHOULD</b> be capable (dependent on configuration) of automatically adapting its congestion response to coexist with TCP Reno congestion controls [ <a href="#">RFC5681</a> ]. To participate in the L4S experiment, a scalable congestion control <b>MUST</b> be capable of being replaced by a Classic congestion control (by application and by administrative control).		As a requirement, such a monitoring being a <b>MUST</b> seems right, but doubtful about the accuracy of such a detection. The draft mentions using mean deviation in TCP's smoothed RTT in addition to other heuristics, but the concern is false positives.
5. Reduce RTT dependence (A1.5)	<a href="#">Section 4.3</a> : A scalable congestion control <b>MUST</b> eliminate RTT bias as much as possible in the range between the minimum likely RTT and typical RTTs expected in the intended deployment scenario.		Again, agreed with the rationale behind this and the <b>MUST</b> compliance. This might be easy to implement as well based on heuristics but will require thorough testing.

6. Scaling down to fractional congestion windows (A1.6)	<a href="#">Section 4.3</a> : A scalable congestion control <b>SHOULD</b> remain responsive to congestion when typical RTTs over the public Internet are significantly smaller because they are no longer inflated by queuing delay.		It is a good starting point for scalable CCs. But we would <b>wait for some test/real deployment data before fully supporting this requirement.</b>
7. Measuring Reordering Tolerance in Time Units (A1.7)	<a href="#">Section 4.3</a> : A scalable congestion control <b>SHOULD</b> detect loss by counting in time-based units, which is scalable, as opposed to counting in units of packets (as in the 3 DupACK rule of <a href="#">RFC 5681</a> TCP), which is not scalable. This requirement does not apply to congestion controls that are solely used in controlled environments where the network introduces hardly any reordering.		<b>Disagree that using time only and not packet count is a foolproof solution.</b> What if the time threshold value cause slow recovery in case of an actual packet loss event? Would it be better to use either packet count or time threshold? We currently don't support RACK - so time based loss detection wouldn't be possible.
(Burst limit)	<a href="#">Section 4.3</a> : A scalable congestion control is expected to limit the queue caused by bursts of packets. It is only required that the specification of a particular scalable congestion control <b>MUST</b> define, quantify and justify its approach to limiting bursts.		<b>Agree with how the draft doesn't specify any requirement and instead provides the rationale why this can be a problem for maintaining low latency.</b> We currently don't support pacing but grow congestion window somewhat conservatively to avoid bursts.
<b>Scalable Transport Protocol Optimizations</b>	<b>Appendix text (no normative refs)</b>		
1. Setting ECT in TCP Control Packets and Retransmissions (A2.1)	This item only concerns TCP and its derivatives (e.g. SCTP), because the original specification of ECN for TCP precluded the use of ECN on control packets and retransmissions. To improve performance, scalable transport protocols ought to enable ECN at the IP layer in TCP control packets (SYN, SYN-ACK, pure ACKs, etc.) and in retransmitted packets. The same is true for derivatives of TCP, e.g. SCTP.		<b>Strongly agree, wouldn't mind if this becomes a requirement.</b> Currently, we only do classic ECN.
2. Faster than Additive Increase (A2.2)	It would improve performance if scalable congestion controls did not limit their congestion window increase to the standard additive increase of 1 SMSS per round trip [ <a href="#">RFC5681</a> ] during congestion avoidance. The same is true for derivatives of TCP congestion control, including similar approaches used for real-time media.		<b>This is definitely an area to explore as scalable congestion controllers should scale for higher rates.</b>
3. Faster Convergence at Flow Start (A2.3)	Particularly when a flow starts, scalable congestion controls need to converge (reach their steady-state share of the capacity) at least as fast as Classic congestion controls and preferably faster. This affects the flow start behaviour of any L4S congestion control derived from a Classic transport that uses TCP slow start, including those for real-time media.		<b>Again, strongly agree with this to allow better throughput for short flows as explained in the draft.</b>