From WarpTCP to WarpRATE

Why we retired a window for a rate

Dennis Vadura, BADU Networks

November 2026

WarpTCP ran in WarpEngine from the beginning. It is CUBIC done carefully: tuned HyStart, a TCP-friendly floor, and fixes stock CUBIC lacked. We have retired it in favor of WarpRATE.

1 Network fabrics are getting faster

Serializing a 1500-byte packet takes 12 µs at 1 Gbit/s, 1.2 µs at 10 Gbit/s and only 120 ns at 100 Gbit/s. The batching in hosts and network links does not reduce at the same rate: a 64 KB TSO train still leaves as one burst, interrupt coalescing holds ACKs for tens of microseconds, and wireless schedulers operate on millisecond time scales. Each step up in wire speed makes that batching span more packets and every ACK-clocked decision gets noisier as a result. This is noise increase is occurring at a time when the bandwidth-delay product the sender may increase and as a result the cost of a wrong control decision has grow. In essence the network accelerated past WarpTCP's ability to make good decisions about how much data it should send — a better solution was needed.

2 Window control no longer works reliably with noisy ACKs

In WarpTCP a congestion window is clocked by received ACKs. Each arriving ACK does several things but arguably the most important is that it releases the next block of data to send. The problem is that each ACK's timing feeds the pacing and growth decisions for the congestion window and a noisy ACK means the growth and reduction follow a more chaotic regime. Chaos is bad for throughput. On today's paths the ACK stream is rarely smooth. Delayed ACKs, GRO at the receiver, Wi-Fi frame aggregation, LTE grant scheduling and Starlink's reconfiguration intervals all hold ACKs back and then deliver them in a batch. The window turns each batch into a burst on the wire, and the controller reads the compressed ACK train as a full buffer and cuts the transmission rate. WarpRATE was designed to avoid this coupling and sends data using a pacing rate controlled by the quality of data delivery. No errors, keep sending, indications of a queue or errors slow the rate. WarpRATE also uses an elastic time buffer to give ample room for variable delayed ACK's to arrive indicating successful delivery.

3 Measuring delivery volume above the quantum is essential

Every NIC and receiver batches traffic, and the result is that ACKs are scheduled in groups. The groups have minimum inter-group time that is a function of the RTT and whatever else is going on at the receiver. We call that interval a measurement quantum. Below the quantum, a timing signal measures the batching and says little about the path. WarpRATE extracts rate as V, the volume delivered over at least eight epochs, a span long enough to integrate path imposed batching, such that the measurement is immune to the quantum by construction.

4 Separating ALLOW, LOSST and RATE leads to a more stable control law

In a congestion window based controller design one number does three jobs. The window bounds the data in flight, it sets the sending rate, and it is used to control loss tolerance. Any signal that touches one job impacts all three, so reordering read as loss cuts the rate, and a jitter allowance in the loss detector inflates the data in flight. WarpRATE gives each job one control. ALLOW bounds the bytes in flight, RATE sets the inter-packet delay, and LOSST decides when an unacknowledged segment is lost. ACK timing reaches only LOSST, where it widens the interval. The pace is impacted only by delivered volume and by distress signals: loss marks, CE echoes and self-inflicted delay. When a jitter spike widens LOSST, the pace and the bytes in flight remain unchanged. This decoupling has been highly effective at protecting throughput during periods where ACK timing varies due to the earlier mentioned causes.

WarpRATE is a modern congestion control algorithm that we believe is beneficial, together with PFQD, in both long distance internet paths and short intra-datacenter paths. Our preliminary real world deployments are encouraging as they confirm our laboratory testing and validation results.