



PD-sep vs Combined headline latency (trace=w600_r0.0015_st30, 850 reqs; per-config N shown under labels)


TTFT p50 (s)



TTFT p90 (s)



TPOT p90 (s)



E2E p50 (s)

