# Live-Class Scaling & Topology Checklist

Pick the right shape for a live class before you build, and price each tier. Companion to the Fora Soft Learn article on scaling the live class.

## A. Count senders, not headcount

- [ ] Count how many people send video/audio at once at your largest real class.
- [ ] Count the watchers separately - they set delivery cost, not architecture.

## B. Pick the topology

- [ ] <=6 and all talk: a peer-to-peer mesh is fine (a tutorial).
- [ ] Dozens who may talk: one SFU media server (a seminar).
- [ ] Few talk, many watch: split into an interactive tier and a broadcast tier.

## C. Turn on simulcast or SVC

- [ ] Confirm the platform sends multiple quality layers (simulcast / SVC).
- [ ] Without it, one weak-network learner sets the quality for the whole class.

## D. Plan the broadcast tier

- [ ] Bridge the SFU to low-latency HLS (LL-HLS) for the silent majority.
- [ ] Deliver via CDN; add a promote-to-speak path for called-on learners.

## E. Capacity and regions

- [ ] One SFU node holds a region (~500-800 senders); cascade across regions.
- [ ] Route each learner to the nearest node to raise the ceiling and cut latency.

## F. Track and caption at scale

- [ ] Define attended / completed up front; emit xAPI on both tiers; batch events.
- [ ] Caption the source once (WCAG 2.1 AA, SC 1.2.4); fan text to both tiers.

## G. Build vs buy

- [ ] Never build the media engine - rent a managed API or run an open-source SFU.
- [ ] Spend engineering on the learning layer: promote-to-speak, tracking, LMS bridge.