

WebRTC Live-Learning Readiness Checklist

Pressure-test a live-learning product against the four classroom gotchas before you pilot. Companion to the For a Soft Learn article on WebRTC for live learning.

A. Latency budget

- Target sub-second glass-to-glass delay for the interactive tier (aim 150-400 ms).
- Confirm the use case is a conversation (WebRTC), not a one-way broadcast (HLS).

B. Network, firewall, and TURN cost

- Test on a locked-down school/corporate network (ports 80/443 only), not just home Wi-Fi.
- Run TURN over TLS on port 443 so relayed media traverses strict firewalls.
- Model the relay-bandwidth bill at your real relayed-learner share before pricing.

C. Scaling topology

- Rule out a peer-to-peer mesh beyond ~6 active video senders.
- Use an SFU for the interactive seminar; add a broadcast tier for the lecture hall.

D. Recording for the catalog

- Record server-side from the SFU - never rely on the instructor's browser tab.
- Trim, compose, transcode, and publish via CDN; emit xAPI Video Profile statements.

E. Captions and compliance

- Ship real-time live captions (WCAG 2.1 AA, SC 1.2.4) - a procurement requirement.
- Confirm WebRTC media is encrypted by default (DTLS/SRTP) for buyer security review.

F. Build vs buy

- Never build the WebRTC media engine - rent a managed API or run an open-source SFU.
- Invest engineering in the learning layer: roles, tracking, recording, LMS bridge.