

1 - Blur your background right now (every app runs on-device)

Zoom	Settings > Background & Effects > Blur (or ^ next to Stop Video)
Zoom mobile	More > Background & Effects (iOS) / Virtual Background > Blur
MS Teams	Background filters > Blur; in call: More > Effects and avatars
Teams shortcut	Ctrl+Shift+P (Windows) / Cmd+Shift+P (Mac)
FaceTime / iPhone	Tap your tile > Portrait; or Control Center > Video Effects
Continuity Camera	Mac Control Center > Video Effects > Portrait
Google Meet	More (3 dots) > Apply visual effects > Blur

2 - Minimum hardware (2026 - re-verify per client release)

Zoom	v5.5.0+, ~4 GB RAM; mobile iPhone 8+ / iPad Pro 5th gen+
MS Teams	CPU with AVX2; not on Linux or virtual desktops (VDI)
iPhone Portrait	A12 Bionic chip+ (iPhone XS+), iOS 15+
Continuity Camera	iPhone XR+; brings Portrait to Intel Macs too

3 - How it works: segmentation vs depth

Step 1 - SEGMENT: a small neural net labels each pixel person/background -> a mask (stencil).
Step 2 - BLUR: keep sharp where mask=person, blur where mask=background, then composite.
Zoom/Teams/Meet = segmentation only (flat image, uniform blur, any webcam).
iPhone Portrait = depth-based (measures distance, graduated bokeh) - looks more natural.

4 - The 33 ms-per-frame budget (why it needs a GPU)

30 fps -> 1 s / 30 = 33.3 ms per frame. Segment + blur must finish inside that window.
MediaPipe Selfie Segmentation: < 3 ms on GPU vs 90-120+ ms on CPU. Same model, different chip.
Model: MobileNetV3-based, 256x256 input, ~106K params, ~454 KB.

5 - Build it into your own product (browser) - checklist

- Tap the raw-frame point: MediaStreamTrackProcessor in, MediaStreamTrackGenerator out.
- Segment with MediaPipe Image Segmenter; run the blur + composite on the GPU (WebGL2 / WebGPU).
- Feather the mask edge a few pixels to kill halos at hair and glasses.
- Blend each mask with the previous frame to stop the outline shimmering (temporal flicker).
- Call frame.close() on every VideoFrame, or the tab leaks memory and crashes.