Face Effects (Beauty / Gaze / AR) - Engineering Cheat Sheet

1 - One map, three jobs (the shared foundation)

Beauty filter	uses the mesh to find skin (smooth) and feature points (reshape)
Gaze correction	crops the eye region from the landmarks, redraws eyes to camera
AR effect	anchors a 3D object using the mesh shape + head-pose matrix

2 - The map: MediaPipe Face Landmarker (free, on-device)

```
478 landmarks  = 468 face-surface points + 10 iris points (gaze uses iris).
52 blendshapes = expression dials 0..1 (jaw-open, brow-raise, smile...).
+ transformation matrix = head position/tilt for placing AR objects.
```

3 - Where it runs in a WebRTC call (Insertable Streams)

```
Camera -> MediaStreamTrackProcessor -> [VideoFrame stream] -> worker
worker: Face Landmarker -> apply effect -> render (WebGPU) -> frame
frame -> VideoTrackGenerator -> encoder -> network.  Run it in a WORKER.
```

4 - The one number: the frame budget

```
30 fps -> 1000 / 30 = 33.3 ms per frame to read + process + return.
Landmarks on GPU ~16 ms (fits). On CPU ~70+ ms (stutters). Use the GPU.
```

5 - The EU AI Act line (do not cross by accident)

```
CHANGING a face (beautify / gaze / AR) = not regulated as risky AI.
READING expression to infer emotion = emotion recognition; banned at work/school.
```

6 - Build-vs-buy checklist

- [ ] Custom effects + on-device privacy + no per-user fee? -> build on MediaPipe Face Landmarker.
- [ ] Need a lens library live next week? -> buy an SDK (Banuba / Snap Camera Kit / DeepAR).
- [ ] Need gaze correction on web/mixed devices? -> NVIDIA Maxine Eye Contact (GPU).
- [ ] Users in FaceTime on a recent iPhone? -> Apple already does eye contact; do nothing.
- [ ] Move per-frame work to a worker on the GPU - never the main thread.
- [ ] Touching expression/emotion? -> get legal review before shipping in the EU.