
Reference architecture for on-device accessibility on smart glasses

Recurring patterns from three Rokid AR Glasses builds

Jennifer Mason / Creative Devs / jennthede.com

1. Two-app architecture.

Glasses are the I/O surface: microphone, HUD, touchpad. Phone holds all ML, all state, all sensitive data. The glasses see nothing they don't need to.

2. No INTERNET permission on the glasses APK.

Verify with `aapt2 dump badging`. Reject any change that re-introduces it.

3. HUD is monochrome.

Severity encodes through weight, italic, brightness, position, flash timing. Not color.

4. Transport channels.

BLE GATT for control: notify-only, MTU 517, distinct service UUIDs per project so multiple apps can coexist on the same glasses.

Wi-Fi P2P for audio: phone is group owner at 192.168.49.1, glasses join with credentials pushed over BLE. 16 kHz mono PCM at roughly 100 ms per frame over TCP.

The Wi-Fi P2P group has no internet capability. Even if a process tries to route over it, traffic dies at the link layer.

5. On-device ML stack.

Speech-to-text: Vosk as primary, Whisper as fallback chain.

Translation: NLLB-200.

Text-to-speech: MMS-TTS.

Audio event classification: YAMNet on TFLite.

6. Confidence gating and hysteresis.

Per-concept thresholds, derived empirically per concept and per environment.

N-tick hysteresis to suppress single-spike false positives. Tier-1 safety concepts get three ticks, which is about 1.5 seconds.

Clipping gate for impulsive false positives from microphone saturation.

7. Data tier model (T0 through T3).

T0, air-gapped: glasses-side state that does not persist.

T1, standard: UI prefs, model files, build metadata. Unencrypted, on the device.

T2, personal: encrypted, opt-in, user-erasable.

T3, sensitive: SQLCipher session DB, ephemeral RAM ring buffer, hardware-backed keystore, Argon2id KDF. Audio is T3 by default for any project doing live capture.

8. Foreground service for sustained audio capture.

`foregroundServiceType="microphone"` plus partial wake lock. Dispatch the pipeline off the main thread so the audio probe cascade doesn't ANR.

9. Backup-deny on both APKs.

`allowBackup=false, dataExtractionRules, fullBackupContent`. No accidental cloud backup of user content.

10. Test floor is mandatory.

Each project has a green-test count from the prior session. Don't ship below it. Don't fabricate test results to hit it. If a test wasn't run on real hardware, it's `@Disabled` with the reason.

Talk: Accessibility-First XR. Seattle XR Connect, June 6, 2026.