HumanityAI'd — Patrol Sight

Your officer looks at a face. They know who it is before the person walks past.

Patrol Sight puts face recognition into a security officer's line of sight. The officer wears lightweight AR glasses; what they see streams to a server **inside your building**, which identifies faces against **your own watchlist** and returns a colour-coded label into the officer's lens in about **68 milliseconds** — hands free, no phone, no radio call, no waiting for the control room.

Why security operations choose it

- **Recognition where the officer is** — not at a gate, not in the control room, not after the encounter.
- **Measured, not claimed** — ~68 ms glasses-to-label (49–103 ms range), 14–20 recognition frames per second, published stage by stage.
- **Your watchlist, your server, your rules** — nearly 94,000 people enrolled in the live reference deployment, with a documented path to 13 million. No scraped database, no cloud, no vendor-held face data.
- **Sovereign by construction** — biometric data never leaves the building, satisfying UAE and Saudi data-residency requirements by architecture rather than by promise.
- **A human always decides** — every match is presented to an officer to verify. The system informs a decision; it never authorises one. Every search and match is logged and attributable.

Proven, not promised

A full production deployment runs today: a **34-service stack on a single GPU**, **93,728 enrolled persons**, glasses-to-label recognition measured at **~68 ms typical** with **0% packet loss**, and a watchlist search P95 of **8.5 ms** across nearly 2,000 real searches. The same software, deployment tooling and operator console you would receive.

How it works

AR glasses stream hardware-encoded H.264 over your dedicated wireless network → **your on-premises GPU server** detects faces and matches them against your watchlist (ArcFace on TensorRT + GPU FAISS) → a **colour-coded category label** appears in the officer's lens while the control room sees the same alert live. A two-layer tracker answers 97.6% of frames from cache, which is what makes continuous recognition affordable on one GPU.

What it does not do — stated up front

Reliable identification to about **5 metres** (6–7 m in good light, 8–10 m in long-range mode) — beyond that, fixed cameras remain the right tool. No automated enforcement action. No emotion or demographic inference. No covert-operation features. We publish the envelope because a security buyer will find it anyway.

Ideal first deployment

A high-footfall public realm or transport environment where officers already patrol on foot and a control room already runs facial recognition on fixed cameras. Start with a **6-week, one-site pilot** — including an information-security and data-protection gate before any live watchlist is loaded — and carry the watchlist, thresholds and audit history straight into rollout.

Let's scope your pilot. · HumanityAI'd · Artificial Intelligence for Smart Glasses