

Demo Script & Talk Track

Patrol Sight — live demonstration

HumanityAI'd — July 2026

CONFIDENTIAL — INTERNAL SALES ENABLEMENT

0. Before the meeting (setup checklist)

- [] Glasses charged and provisioned to the demo server; server on a **5 GHz** demo network with power-save disabled and a static IP. *(Test the ping. If it is not single-digit milliseconds, fix the radio before you fix anything else.)*
- [] Demo watchlist loaded with **consenting colleagues only**, each assigned a category and colour. Never demonstrate on a real watchlist, and never on a person who has not consented.
- [] Verify TensorRT is active, not the ONNX fallback — check the engine loaded cleanly at startup.
- [] Command console open in a browser tab: live device monitor, alert stream, audit log.
- [] Run one live recognition end to end before the client walks in.
- [] Have printed face photos of the consenting cohort as a backup recognition target if the room is badly lit.
- [] Know your numbers cold: ~68 ms typical, 93,728 enrolled, P95 search 8.5 ms, 97.6% cache hit, reliable to 5 m.

1. Open (2 min) — the gap, not the technology

"You already run facial recognition. Your control room has it on thousands of cameras and it works. So here's the only question I want to ask: when one of your officers is standing three metres from someone who's on your watchlist — does the officer know? Or does the control room find out afterwards, and radio it through once the person has already walked on?

That gap — between what your cameras know and what your officer knows, in the moment — is what we close."

Do not open with AI, models, or architecture. Open with the operational gap. If they push back that the control room does relay in real time, ask how long it takes and whether the officer has hands free to answer.

2. The 'wow' — live recognition in the lens (4 min)

1. Hand the glasses to the client and have them put them on. Let *them* wear it — this lands far harder than watching you wear it.
2. A consenting colleague walks into their field of view. **The label appears in the lens as they look** — name and a category colour.
3. Have a second colleague enter, with a different category and colour. Both labels track independently and stay stable as heads move.

4. Ask the client to turn away and look back. The label reacquires.

"That was about seventy milliseconds — glasses to lens. You didn't press anything, and it wasn't a photo you took and submitted; it's continuous. And that match ran on the box in this room, against a list on that box. Nothing went to a cloud, because there is no cloud in this product."

Then show the control room side. Switch to the console: the same alert with device, timestamp and confidence, and the live device monitor showing frame rate and latency per pair of glasses.

"Your officer sees it and your operations centre sees it, at the same time, from the same event."

3. The scale proof (2 min)

Open the enrolled population count.

"Ninety-three thousand seven hundred and twenty-eight people in this index, on one mid-range GPU. Search comes back in eight and a half milliseconds at the ninety-fifth percentile — we measured that over nearly two thousand real searches, not a benchmark script. We've measured the same pipeline out to thirteen million faces: eighty to a hundred and thirty milliseconds. Your watchlist size is not the thing that will limit you."

If they are technical, give them the reason: **the two-layer tracker skips the vector search on 97.6% of frames.** That single design decision is why this runs on one GPU instead of a rack.

4. The watchlist story — their control, not ours (4 min)

Open the console.

1. Enrol a new consenting person live: create the record, upload a photo, assign category and colour.
2. Show the index rebuild and hot-reload — "no restart, no retraining, no maintenance window."
3. Walk that person into the officer's view. They are recognised with the new category colour.
4. Then open the **audit log**.

"Two things to notice. First, your team did that in under two minutes without calling us. Bulk enrolment runs at forty to sixty images a second, so onboarding an existing list is a morning's work, not a project.

Second — and this is the part that matters when your oversight body asks — every one of those searches is logged and attributable. Who searched, which device, when, what confidence. We don't supply you with a face database; you enrol who you're lawfully entitled to enrol, and the record proves it."

5. The limits — say them before they ask (3 min)

This section wins the technical evaluation. Do not skip it to protect the mood.

"Let me tell you what it doesn't do, because your engineers will find this out anyway and I'd rather you hear it from me.

Range. Reliable to about five metres. Six or seven in good light. There's a long-range mode that reaches eight to ten at a lower frame rate. Beyond that a head-worn camera doesn't have the pixels, and we reject the face rather than guess. If someone tells you their glasses identify people across a plaza, ask what happens at the pixel level.

The network matters more than the AI. On a 2.4 gigahertz link with power-save on, our round trip was about a hundred milliseconds — bigger than the whole recognition pipeline. On 5 gigahertz WiFi 6 with power-save off, four. So a dedicated wireless network is a requirement in our deployment guide, not a nice-to-have.

We don't have a NIST ranking. NEC and IDEMIA do, and they'll tell you so. What we have is a fully measured end-to-end pipeline in a live deployment, and we'll measure accuracy in *your* environment during the pilot and contract on those numbers. Ask the NIST-ranked vendors what their glasses-to-lens latency is.

We have security hardening still to finish. Credential rotation, per-device tokens, encrypting the wearable link. It's written up, costed, and dated in our due-diligence report, which I'll leave with you. Until the link is encrypted, we specify an isolated network for the glasses — a control, not an omission."

"And one thing it will never do: act on its own. Every match goes to an officer to verify. It informs a decision; it doesn't make one."

6. The fit — why on-premises wins here (2 min)

- Biometric data is sensitive personal data under UAE and Saudi law, with in-country processing the default. On-premises satisfies that by architecture, not by contractual assurance.
- The watchlist, the faceprints, the alerts and the audit trail are all on their server. It runs air-gapped if they want it to.
- Neutral supply chain — not cloud-dependent, and not exposed to the sanctions position of the largest surveillance vendors.
- One server per site, deployed and recovered by their own IT with one command.

7. Objection handling

Objection	Response
"NEC / IDEMIA already do our biometrics."	They do gates and cameras — fixed points the subject walks to. We do the officer who walks. We deploy alongside them; we're not asking you to replace anything.
"What's your accuracy, independently verified?"	No NIST ranking today, and I won't imply one. Here's our internal 0-90° pose benchmark, and here's the offer: we measure accuracy in your environment during the pilot, including the false-positive rate, and contract on what we measure.
"Chinese vendors sell glasses cheaper."	They do, and they're on the US Entity List with no GCC presence or support and no PDPL story. On a project with international exposure that's a procurement problem, not a price comparison.
"We tried smart glasses here years ago."	You did — a fifty-unit deployment in UAE security operations in 2019, and it worked. That's my best evidence. That was 2019 hardware with an offline list; we bring current glasses, a measured sub-100 ms pipeline, a 94,000-person GPU index and a full command console.
"Facial recognition is a legal and PR risk."	The risk lives in scraped databases, cloud processing and automated action — we have none of the three. Your list, your server, your operators, human verification mandatory, everything logged. We'll support your DPIA in the pilot.
"How many officers can one server handle?"	I don't have a measured number yet, so I won't invent one. It's bounded by video decode and detection, not by search. We measure it in your pilot and price the fleet against that figure.
"Officer comfort over a twelve-hour shift?"	Lightweight AR glasses, and we run a comfort survey in the pilot with rotation and spare units. There's also a category-only display mode if you don't want names shown to the wearer.
"What if the wireless drops?"	The session survives brief interruption and reacquires. Sustained loss means no recognition — the officer is back to normal patrol, not looking at stale data. Degraded on-device matching is on the roadmap.
"Price?"	The pilot is a fixed, contained fee. Full pricing scales with watchlist size and fleet — let's scope it against your actual site.

8. Close (1 min)

"Let's put this on five officers at one of your sites for six weeks. Your patrol area, your watchlist, your wireless network, measured against numbers we agree today — including the false-positive rate, which we'll establish in a shadow-running week before anyone acts on a match. Your information-security team gates it before a single live record is loaded.

If it doesn't change what your officers can do, you've risked six weeks and no operational commitment. Which site would you want to try it in?"

Next step: agree the pilot site and start date, book the RF survey, name the information-security reviewer and the watchlist owner.