

Facial AI Verification Platform

Consent-based, one-to-one identity verification for controlled access and service workflows

1:1

CLAIMED IDENTITY
VERIFICATION

Non-technical overview
and technical design

PART 1

What the app does

01

A plain-language view for users, decision-makers and demonstration audiences

A safer way to confirm a claimed identity

The person states who they are, then the app compares their live face only with that enrolled record.

Controlled entry

Doors, secure cabinets and restricted rooms

Workstation access

Local confirmation before protected work

Service verification

Identity confirmation in an approved workflow

Equipment release

Short-lived authorization for controlled assets

The prototype does not search crowds or identify unknown people.

The user journey

Enrollment records consent first. Verification then checks a claimed identity in a guided live workflow.

1

Sign in

An operator uses an authenticated session.

2

Enroll

The volunteer provides consent and 8 to 15 guided views.

3


Verify

The person enters their identifier and completes live face checks.

4

Decide

The app returns a clear result and records an audit event.



Captured images support the immediate operation. The database stores encrypted face templates rather than a photo gallery.

Privacy and user control

The design limits collection and provides a direct way to remove an enrolled record.

Consent before capture

Enrollment requires a stated purpose and an affirmative consent record.

One-to-one comparison

Verification checks only the identifier that the user claims.

Delete saved users

The operator can remove identity metadata, consent and biometric templates.

Fail-closed access

A missing user, incompatible template or failed match keeps access locked.

Government use requires legal review, clear retention rules, human oversight and a non-biometric fallback.

Prototype status

The application works as a local Windows demonstration, but it has not completed production biometric validation.

Available now

- Windows executable
- Webcam enrollment and verification
- Saved-user database and deletion
- Technical and simple result views
- Simulated access points and audit records

Required before a government pilot

- Independent accuracy and bias evaluation
- Validated presentation-attack detection
- Role-based operator controls
- Formal security and privacy assessment
- Approved hardware integration and operating procedures

No production accuracy claim is made in this presentation.

PART 2


Technical design

02

Architecture, biometric pipeline, data controls, deployment and remaining engineering work

Application architecture


The packaged Windows app runs a local FastAPI service and serves the Next.js dashboard on the loopback interface.



The Docker design also supports PostgreSQL, Redis, MinIO and a reverse proxy for a future managed deployment.

Biometric processing pipeline

Every accepted frame must contain one clear face and pass image-quality checks before comparison.



Enrollment averages 8 to 15 normalized embeddings. Verification uses a packaged threshold of 0.42 and checks only the claimed user's latest compatible template.

MODEL: opencv-sface-2021dec-v1

DETECTOR: YuNet 2023mar

Data model and security controls

Identity metadata, consent, biometric templates and event history use separate records with explicit deletion behavior.

User	External identifier, display name, active status
Consent	Purpose, notice version, grant time
BiometricTemplate	Encrypted vector, dimensions, model version
VerificationEvent	Outcome, similarity, latency, model version
AuditEvent	Actor, action, target, detail, timestamp

Fernet protects template bytes at rest. JWT protects operator endpoints. Deletion cascades to consent and templates while preserving a minimal audit record.

Accuracy and anti-spoofing

The current build raises the match threshold and combines identity scores with a three-view head-movement check.

Identity decision

Median score across front, left and right views must reach 0.42.
Every view must also remain within 0.05 of the threshold.

Motion challenge

YuNet landmarks estimate horizontal head-pose variation. The challenge fails when the three captures do not show enough movement.

Security boundary

A moving printed image or screen replay may still defeat this basic challenge. Production use requires a validated PAD model or depth/infrared sensor, calibrated thresholds and independent attack testing.

Higher thresholds can reduce false accepts while increasing false rejects. Evaluation must measure both outcomes on representative, subject-disjoint data.

Deployment, verification and next steps

The project packages the frontend, API, face models and runtime into a Windows x64 application.

Build	Next.js static export + PyInstaller one-directory bundle
Runtime	127.0.0.1:8765 with browser-based webcam access
Persistence	SQLite and installation secret under Local AppData
Verification	8 API tests, frontend TypeScript build, packaged smoke test

Recommended next phase

1. Run a labelled accuracy and demographic evaluation
2. Integrate certified presentation-attack detection
3. Add OIDC, operator roles and rate limiting
4. Complete legal, security and privacy assessments
5. Pilot with approved sensors and human fallback

Current classification: functional technical prototype, not approved for autonomous government decisions or physical access.