

— THE PERCEPTION & TARGETING LAYER FOR ANY DRONE

Cue any target. Hold the track. Hand it off.



The foundation AI tracker that locks onto any target, whether trained, described in text, or shown once. Holds track through motion, lighting changes, and occlusion. Shares that target memory across a swarm to keep custody through handoff. Runs on any drone, any camera. No retraining. All autonomous, in denied environments.

YOU BRING THE PLATFORM. WE BRING THE EDGE AI

The perception and targeting layer that runs under your existing system — piloted or autonomous, any drone, any camera, any mission. Field validated performance.

Find

Acquire the target

Cue any target, however you find it

A click, a class, a plain description, or an example photo. Starts tracking from a single detection, no training dataset required.

See in EO or IR

One model, both spectrums, day or night, no retraining to switch.

Fix

Hold custody

Hold the track through anything

Builds a running picture of the target's appearance, staying locked through frame drops, motion, lighting, and aspect changes, even at low frame rate and on a slow or jittery camera. Adapts online, on lower-quality frames.

Track multiple targets at once

Independent tracks, IDs, and histories across several objects in the same scene.

Finish

Deliver/hand off

Follow all the way through terminal guidance

As the target gets bigger in the final approach and its appearance changes fast, the memory updates automatically. No near- and far-range training data needed.

Hand targets across platforms

Shares the tracker's memory, not just position, so another platform can reacquire after handoff.

Runs onboard the hardware you already fly—50 Hz on Orin NX, 30 Hz on Snapdragon 8, and no high-frame-rate camera required. Built for DDIL conditions.

INTEGRATE INTO THE STACK YOU ALREADY HAVE

Modular tracking layer that drops onto any detector, VLM, or manual cueing workflow already in the pipeline. Cue targets by click, text prompt, image example, or existing detector/VLM output.

Your sensor

EO/IR video via standard RTSP



Latent ATR

Foundation-model tracking + memory



Your platform

Targeting, ISR, autonomy stack

See it lock on your target. Request the ATR integration brief: info@latentai.com

TWO WAYS TO RUN IT

One capability, two ways to deploy it. Meet partners where they are today.

Pilot + GCS

ATR on the operator's screen


The operator confirms the target before anything downstream acts on it. No autonomy required; runs inside your existing flight workflow.

Autonomous / Swarm

Tracks pass directly to your autonomy stack

Multiple targets held across the formation. Scales as your autonomy matures.

WHY IT BEATS THE ALTERNATIVES



Unlike platform-bundled ATR solutions that are tied to one airframe or one sensor, Latent AI's ATR runs under any stack.

LATENT AI VS CONVENTIONAL TRACKER

6.6x higher Success

11.3x higher Weather robustness

3.9x lower CEP

43.5x smaller Minimum pixel size

2.9x farther Distance tracked

9.5x longer Time tracked

YOLO only — Detections only, no persistent track; needs a large labeled dataset per class; comparable accuracy runs ~10 Hz

Pixel lock — Breaks under motion or lighting change; needs a high frame rate just to avoid losing track.

Latent ATR — Persistent track with memory; cues from one detection or a text/image prompt. Holds a fast target even at low frame rate and through platform motion. Adapts online.

PROGRAM CREDENTIALS

USSOCOM CRADA

Army GVSP CRADA

NPS CRADA (Project ODIN)

DIU C-UAS Close-in Kinetic Defeat

CDAO Swarm Forge Crucible 3

Army Operation Jailbreak

AWARDABLE

DIU Marketplace

Tradewinds Solutions Marketplace

Carahsoft

COMPANY CREDENTIALS

DOD IL6 approved

FCL - TS Level

ITAR Compliant

See it lock on your target

Request the ATR integration brief — info@latentai.com

