

PlatePatrol

B2: Christine Li, Vicky Liu, Andy Zhao
18-500 Capstone Design, Spring 2025
Electrical and Computer Engineering Department
Carnegie Mellon University

Product Pitch

Over 850,000 vehicles are stolen in the U.S. every year—that's nearly 2,400 per day. And that doesn't even include vehicles involved in hit-and-runs or AMBER Alerts. Expecting human drivers to identify and report these vehicles in real time is unrealistic. That's why automatic license plate recognition (ALPR) systems are essential. However, most existing systems are fixed in place, offering limited coverage and delayed alerts.

PlatePatrol changes that. We turn everyday dash cams into a real-time, crowdsourced ALPR network. Any driver can purchase and install our smart dash cam, which scans license plates, checks them against a central watchlist, and instantly uploads flagged images—with GPS coordinates and timestamps—to a secure backend. On the other end, law enforcement and third-party services can dynamically subscribe to specific plates by registering their own webhooks—receiving real-time alerts straight into their existing systems.

Built for true end-to-end integration, our dash cam plugs into a car's cigarette lighter and runs lightweight machine learning models on-device for low-latency performance. It scales effortlessly through a distributed cloud architecture, with end-to-end encryption and an instant opt-in/out switch to protect user privacy and system security. Thanks to our API-first design, PlatePatrol integrates seamlessly with existing watchlist systems and tip line services, leveraging what already works without reinventing the wheel. In a 30-minute road test, **PlatePatrol outperformed human observation by 5x**—detecting every plate the driver saw, plus 117 more they didn't.

System Architecture




Fig. 1: Overall System Block Diagram




Fig. 2: Dash Cam Submodules

System Evaluation



Fig. 4: Dash Cam Deployed in Road Test



Fig. 5: Inference Frames

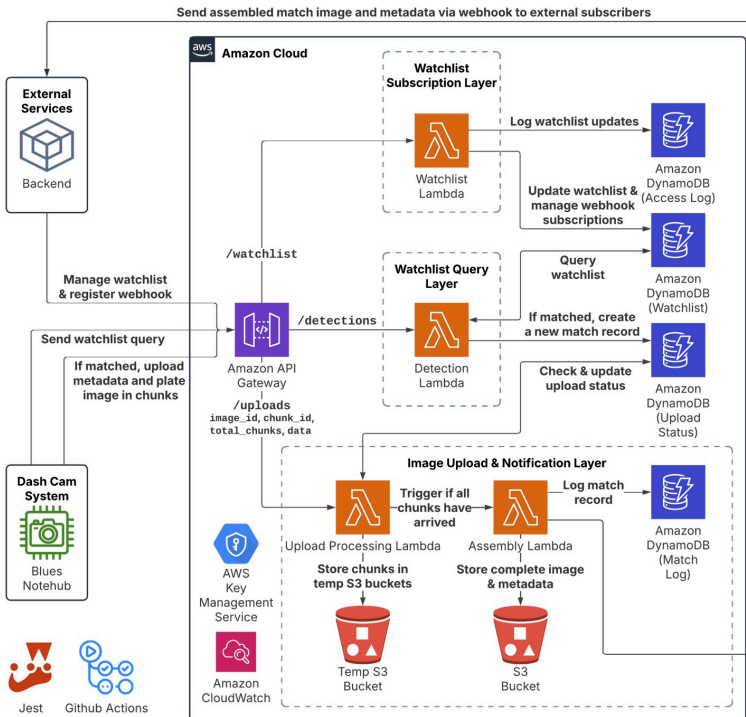


Fig. 3: Central Server Block Diagram

Dash Cam Requirements:

Metric	Target	Result
Basic Functionality	Car-powered 1-minute recording storage	Fully compliant
Ease-Of-Use	≤ 10 minutes installation < 30s startup < 30s shutdown	5 minutes installation 41s startup 332ms shutdown
Privacy	≤ 1s opt-in/out	Instantaneous opt-in/out

ALPR Requirements:

Metric	Target	Result
Plate Detection	≥ 90% mAP50	90.4% mAP50
Text Recognition	≥ 90% accuracy	93.2% accuracy

End-To-End Requirements:

Metric	Target	Result
Accuracy	≥ 80%	100% accuracy + 468% extra plates
Latency	< 1.1s	8.06s
Security	Data encrypted User access controlled	Fully compliant

Design Trade-Off Highlight:

Approach			ALPR Latency	Network Load	Concerns
Plate Detection	Text Recognition	Watchlist Query			
Cloud	Cloud	Cloud	20ms	Frequent full image upload (470s)	Network bandwidth
Edge	Cloud	Cloud	129ms	Frequent plate image upload (7.3s)	Network bandwidth
Edge	Edge	Edge	257ms	Moderate watchlist sync (2.16s) Rare plate image upload (7.3s)	Security & Sync
Edge	Edge	Cloud	257ms	Frequent watchlist query (469ms) Rare plate image upload (7.3s)	/

