



From Roadside to Real-Time: The Journey of an AI Vision Platform

Engineering an End-to-End System for
Automated Container ID Recognition

The Challenge: Manual Asset Tracking at Scale is Broken

Manual tracking of container IDs is a critical but inefficient part of modern logistics. The process is slow, prone to human error, and lacks real-time visibility, leading to costly delays and data inaccuracies. How can we automate this process reliably in complex, real-world environments?



Slow & Inefficient



Error-Prone



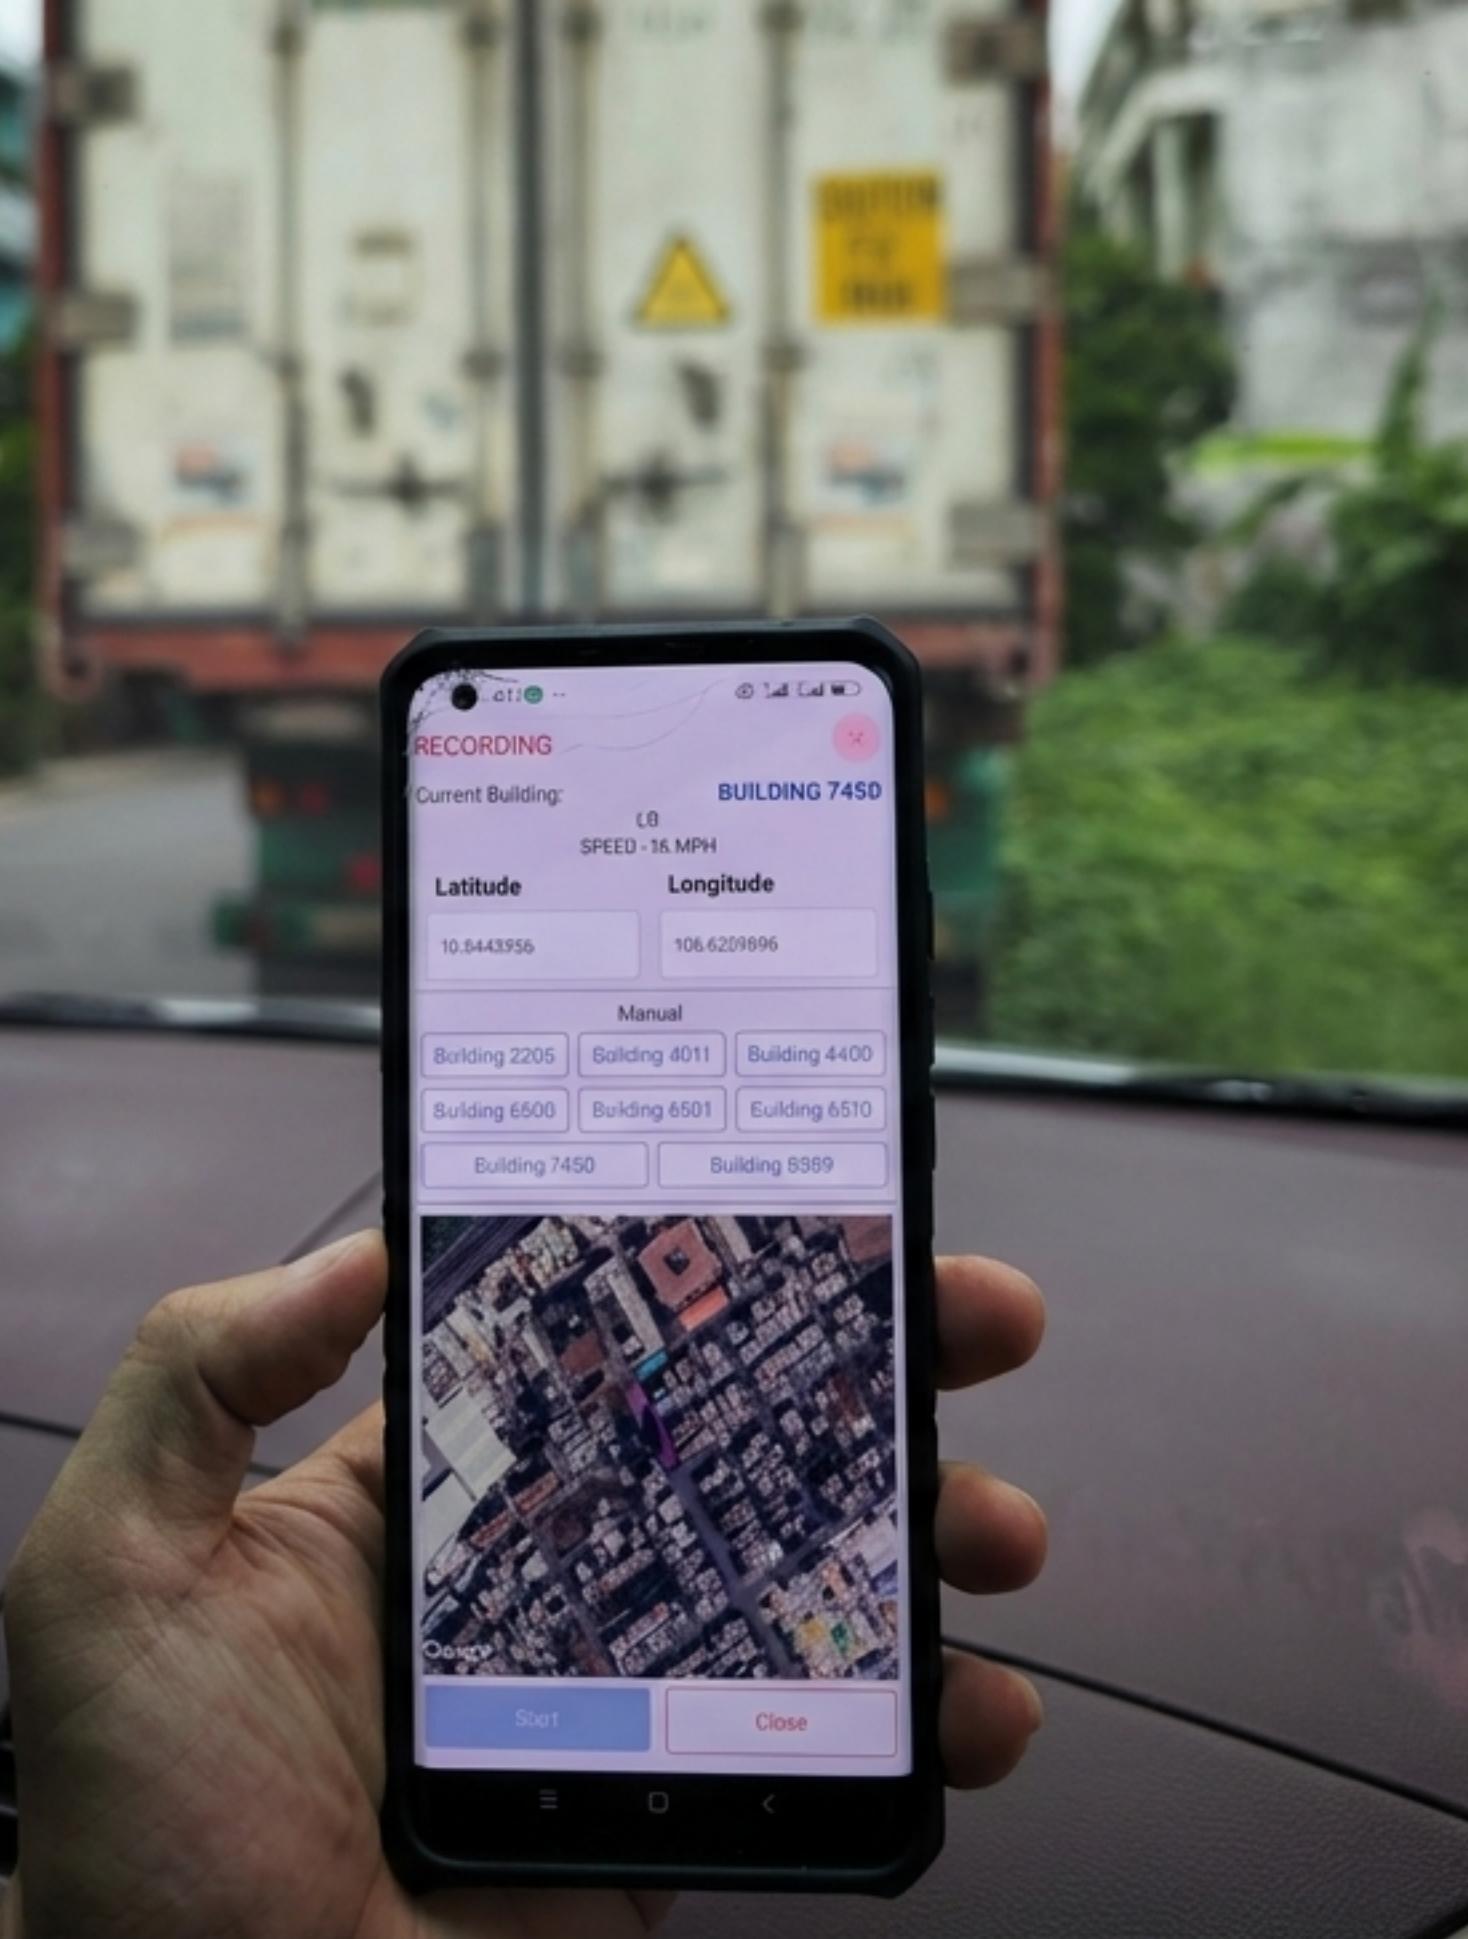
Lack of Real-Time Data



In the Field: Mobile Hardware for High-Quality Data Capture

The foundation of any AI system is its data. We engineered a mobile data collection unit to capture high-resolution video and precise metadata directly from the field. This rugged setup is designed for reliability and performance on the move.



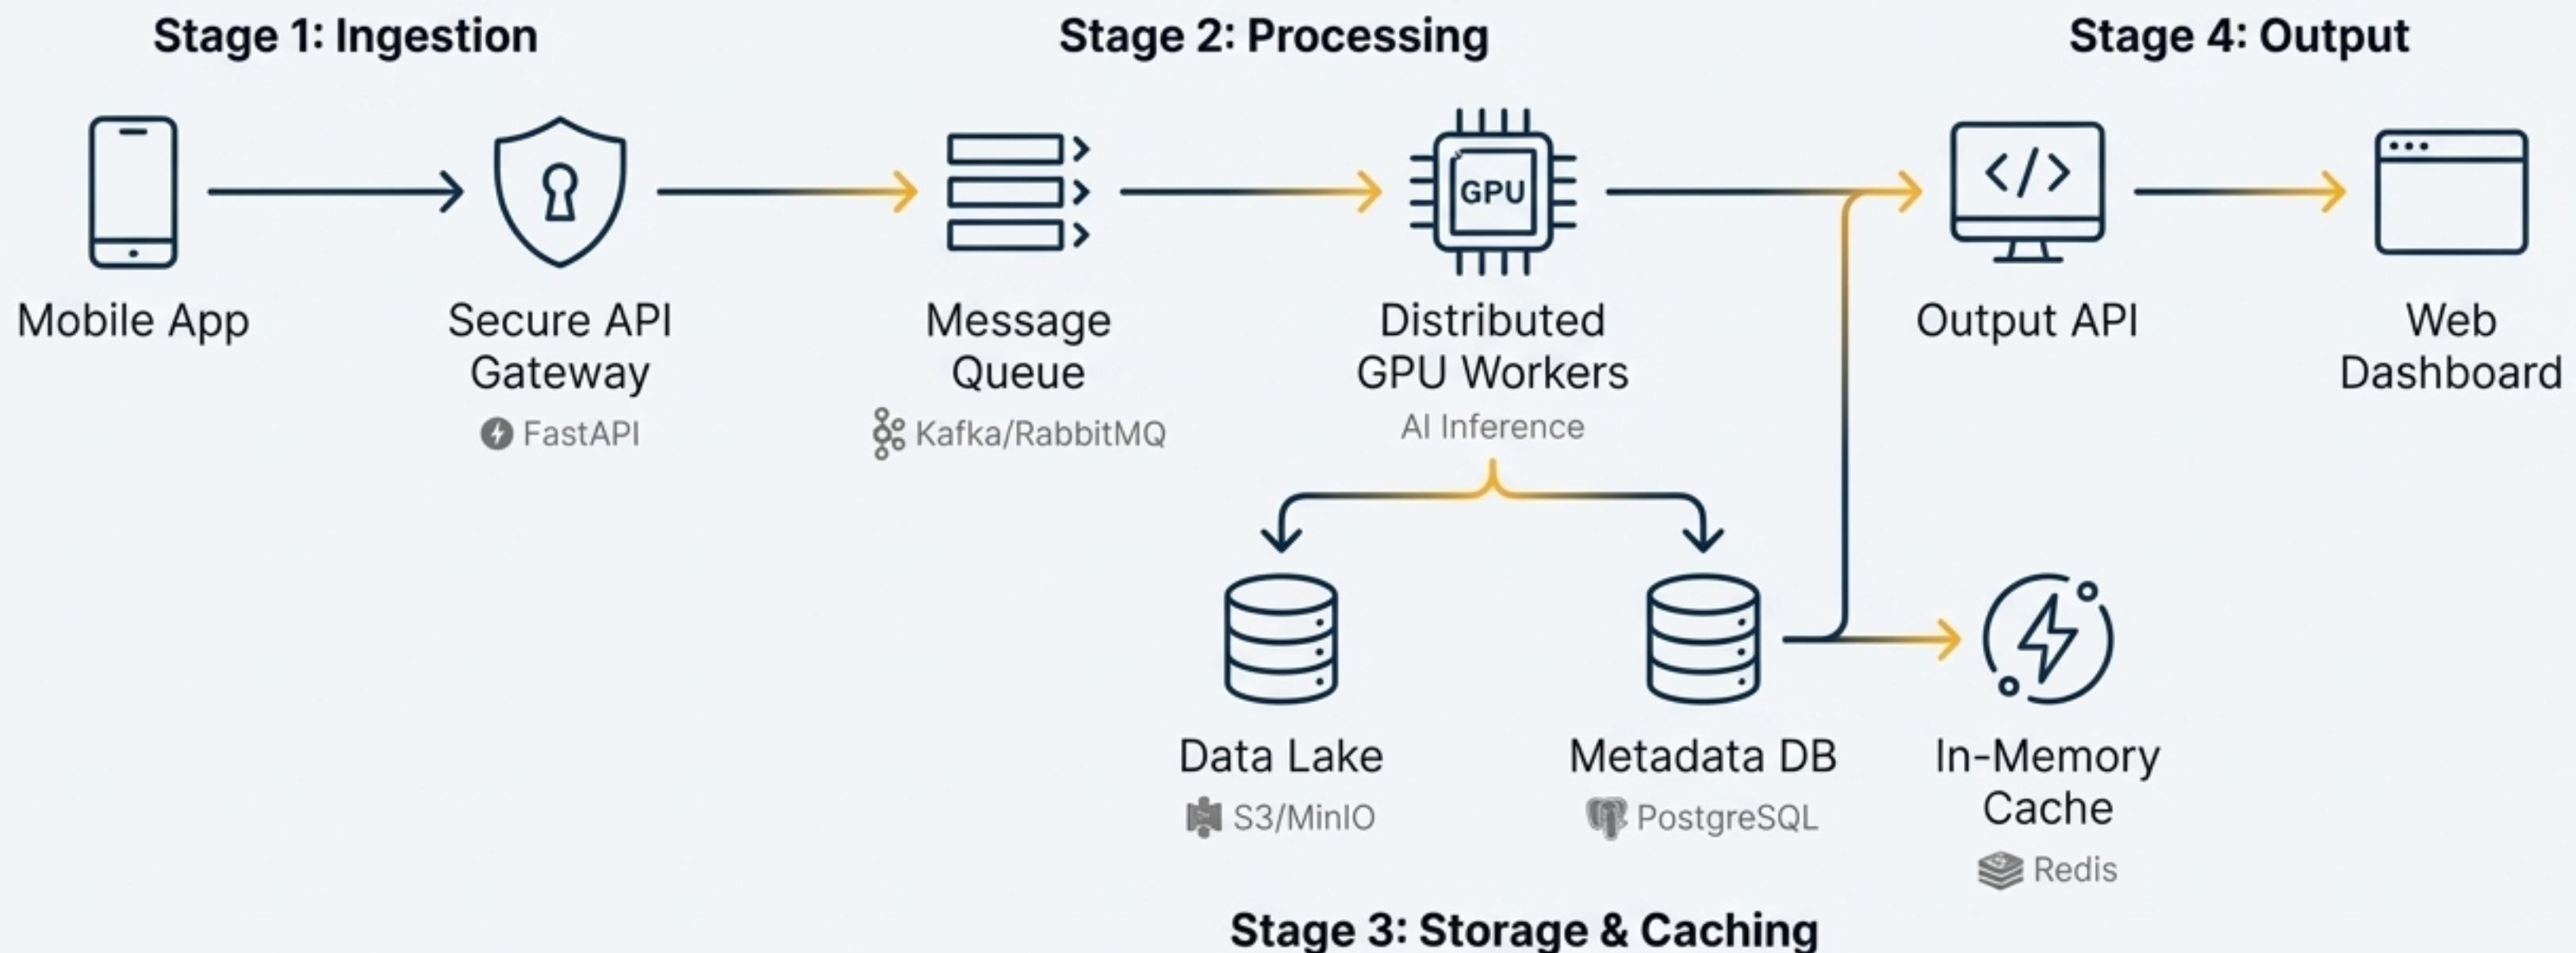


The Mobile Nerve Center: Capturing Rich, Contextual Data

The custom mobile application serves as the primary interface for data collection. It seamlessly fuses the video stream with essential metadata, ensuring every frame is geo-tagged and timestamped for precise backend processing.

-  Real-time GPS coordinate and speed logging.
-  Live map view for operational awareness.
-  Simple interface for manual building/location tagging.
-  One-touch data synchronization to the backend.

The AI Vision Engine: A Scalable Microservices Architecture



Inside the AI Pipeline: From Pixels to Structured Data

Each video frame undergoes a multi-stage process to detect, track, and extract container information with high precision.



Engineered for Speed and Scale

<200ms

Per-Frame Inference
Latency on GPU

Thousands

Images Processed Per Day

System performance is optimized using TensorRT and ONNX Runtime, including model quantization for efficient deployment on GPU servers and edge devices.

Closing the Loop: The Exception Handling Dashboard

No AI is perfect. Our system automates the verification process by flagging low-confidence detections or validation failures for human review. This web-based dashboard is the central hub for quality control and data correction.

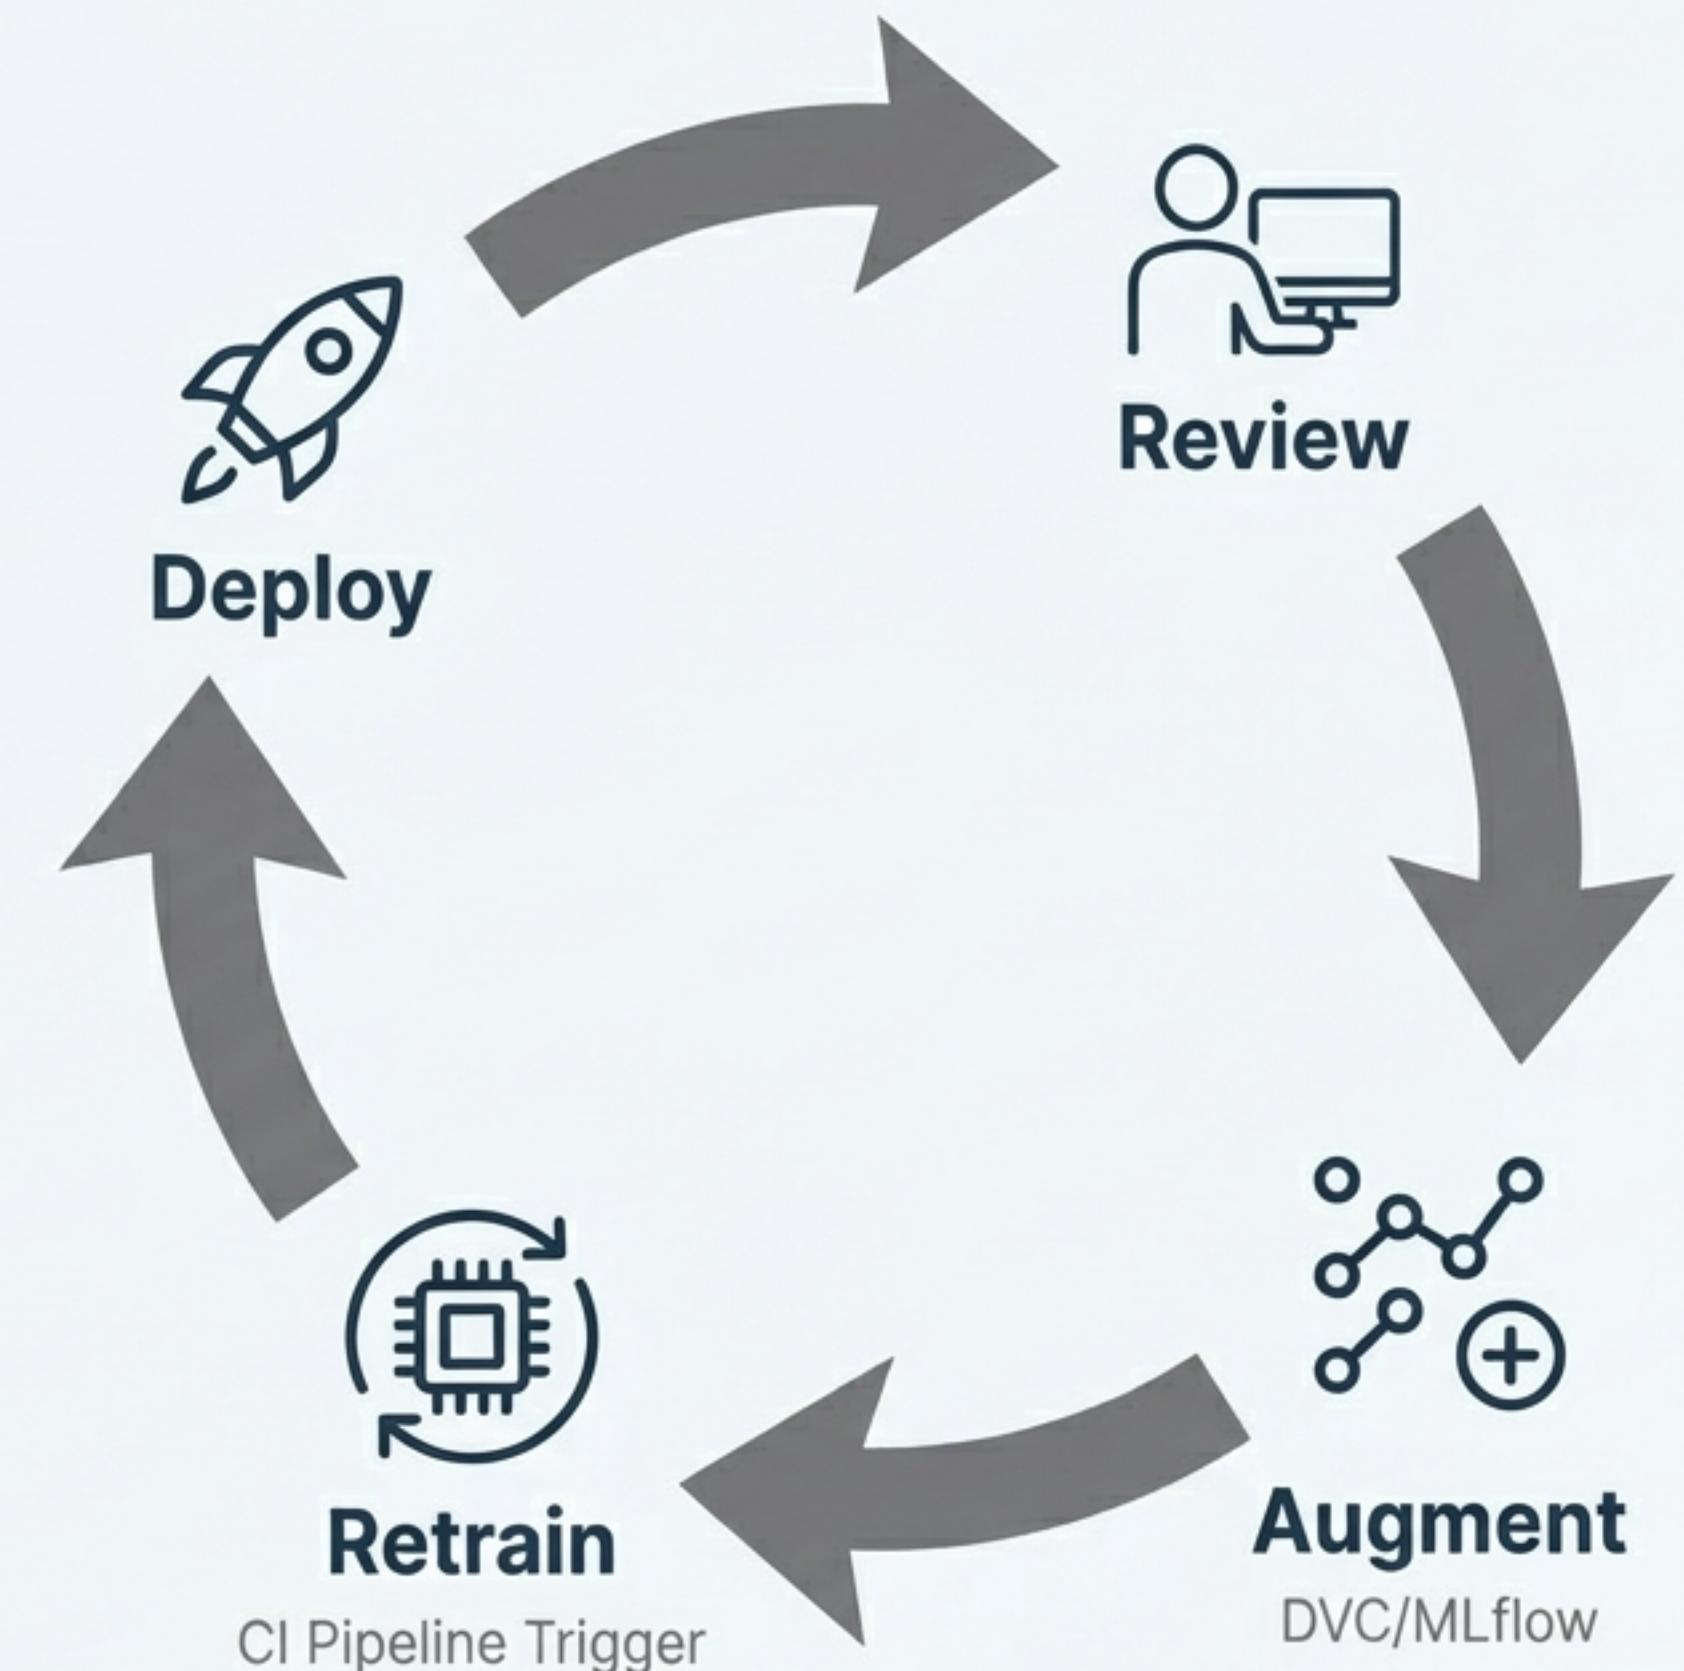
From Date To Date Search Text

Detect Date Time	All Text	Latitude/Longitude	Image	Exception Type	Camera Name	Data Type	Updated By	Updated Date Time
Jul 8, 2024, 3:13:45 PM	CIPUS109477	10.844458, 106.6209063		Missing Container Number or/and Door Number.	1	AI Detect	AI System	Jul 8, 2024, 3:13:45 PM
Jul 8, 2024, 3:12:44 PM	MORU1124831	10.844558, 106.6209063		Missing Container Number or/and Door Number.	1	AI Detect	AI System	Jul 8, 2024, 3:12:44 PM

AI-detected text
Image snippet for review
Latitude/Longitude for context

A Self-Improving System: The Automated Training Pipeline

Human-verified corrections from the dashboard are automatically fed back into our training pipeline. This CI/CD-driven workflow enables rapid model iteration and continuous performance improvement.



Delivering Quantifiable Business Impact



>90%

End-to-end container ID extraction accuracy achieved through iterative training and a robust human-in-the-loop process.



~60%

Reduction in manual verification workload, freeing up operational staff by automating exception detection.



<1 Hour

Model update cycle shortened from days to under an hour, enabling rapid response to changing data and requirements.

Platform Technical Highlights



Secure API Gateway

JWT-based authentication, rate limiting, and TLS enforcement.



Production Model Registry

Full support for versioning, staged deployments, and instant rollbacks.



OTA Model Updates

Over-the-air deployment of optimized models to edge devices.



Comprehensive Observability

Built-in monitoring for inference analytics, accuracy tracking, and data-drift detection.



Distributed Processing

Architected to scale horizontally for increasing data loads.

The Full Journey: A Unified Platform from Capture to Impact

