


SMART INDIA HACKATHON 2026



TITLE PAGE

ID

PROBLEM STATEMENT ID

26187

PS

PROBLEM STATEMENT TITLE

AI-Based Intelligent Video Analytics Platform for Border Surveillance using existing CCTV Infrastructure

TH

THEME

Blockchain & Cybersecurity

CT


PS CATEGORY

Software

TM

TEAM

E2C | Team ID: NA



IBVAP – Intelligent Border Video Analytics Platform

PROPOSED SOLUTION

An Edge-to-Cloud Intelligent Border Video Analytics Platform (IBVAP) powered by NVIDIA's LocateAnything — a state-of-the-art Vision-Language Model that reads existing legacy CCTV streams in real time via natural-language prompts.

HOW IT ADDRESSES THE PROBLEM

1

Hardware-Agnostic

Retrofits existing IP cameras with software-defined intelligence — no expensive proprietary smart cameras needed.

2

Automated Monitoring

Eliminates human fatigue with 24/7 multi-object tracking, intrusion detection & anomaly detection.

INNOVATION & UNIQUENESS

3

Unified VLM Architecture

Replaces multi-model pipelines (YOLO, OCR, FaceRec) with one foundation model — zero-shot grounding, ANPR & activity detection.

4

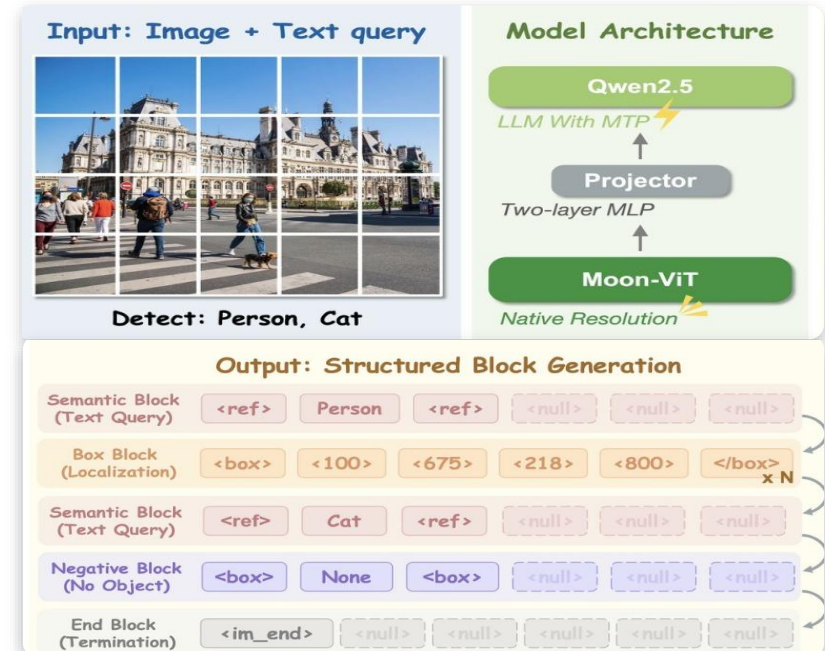
Parallel Box Decoding

Breaks the sequential decoding bottleneck for high-throughput, real-time analytics on edge hardware.

12.7

boxes decoded per second — real-time throughput on edge hardware via Parallel Box Decoding (PBD)

REFERENCE ARCHITECTURE



MODEL HIGHLIGHTS

- Zero-shot grounding
- Built-in ANPR & activity detection
- Edge-optimized inference (TensorRT / ONNX)

TECHNICAL APPROACH

TECHNOLOGIES

AI

CV & AI/ML

NVIDIA LocateAnything (unified detection) · ByteTrack · OpenCV

ST

Streaming

GStreamer · FFmpeg · RTSP · Kafka

BE

Backend & Edge

Python · FastAPI · TensorRT / ONNX Runtime

FE

Frontend & Storage

React.js · TailwindCSS · Leaflet (GIS) · PostgreSQL · Milvus

METHODOLOGY

01

Ingestion

Capture H.264/H.265 RTSP streams from existing BOP IP cameras via GStreamer.



02

Inference

Run frames through LocateAnything with prompts — “Locate humans,” “Read plate.”



03

Tracking & Geofencing

ByteTrack builds trajectories; virtual tripwires trigger on polygon breaches.



04

Output & Alerts

Real-time alerts & snapshot evidence sent to the CCC dashboard via REST/WebSockets.



FEASIBILITY AND VIABILITY



FEASIBILITY

- LocateAnything's Hybrid Mode gives high-precision localization with fast throughput for real-time edge deployment on mini-PCs.
- Integrates with existing CCTV networks via ONVIF/RTSP — no hardware overhaul needed.



CHALLENGES & RISKS

- Running heavy Vision-Language Models on resource-constrained edge nodes (e.g. local BOP servers).
- Accuracy can degrade in severe weather or near-total darkness.



STRATEGIES

- NVIDIA TensorRT with INT8 quantization compresses & accelerates LocateAnything for edge devices (e.g. Jetson Orin).
- Low-light image enhancement preprocessing + thermal/IR camera feeds where available.



IMPACT AND BENEFITS

POTENTIAL IMPACT ON THE TARGET AUDIENCE

Gives the Sashastra Seema Bal (SSB) and other border forces an automated, 24/7 intelligent situational-awareness system — drastically improving mission readiness.

BENEFITS OF THE SOLUTION

OE

Operational Excellence

Zero-latency breach alerts & visual evidence — eliminating blind spots and human observation errors.

ES

Economic Scale

Massively reduces CAPEX & OPEX by reusing legacy infrastructure and consolidating analytics into one stack.

MF

Mission Flexibility

Prompt-based detection lets commanders add new rules instantly (e.g. “Locate cars”) without retraining.

24/7

SITUATIONAL AWARENESS

1 Stack

REPLACES MANY POINT TOOLS

0

RETRAINING REQUIRED

MT

Mathematical Theory & Models

"LocateAnything: Fast and High-Quality Vision-Language Grounding with Parallel Box Decoding" (NVIDIA, 2026).

OG

Operational Guidelines

Ministry of Home Affairs (MHA) border surveillance frameworks and C4ISR/C2 integration standards.

AL

Algorithms

Parallel Box Decoding (PBD), Multi-Token Prediction (MTP) for boxes, DeepSORT / ByteTrack for temporal association.

DS

Data Schemas

RTSP/ONVIF streaming protocols and standard H.264/H.265 video encoding schemas.

NEXT-TOKEN VS. PARALLEL BOX DECODING

Next Token Prediction:

<ref> → ship → </ref> → <box> → <130> → <647> → <911> → <832>

Generic Multi-Token Prediction:

<ref> ship </ref> <box> <130> <647> <911> <832> </box> <box> <122> <233> <451> <677> </box> <ref> cat 399

! Unstructured

Parallel Box Decoding:

<ref> ship </ref> null null null <box> <130> <647> <911> <832> </box> <box> <122> <233> <451> <677> </box>

Box-Aligned



PROJECT REPOSITORY

<https://gitlab.com/lenluarun/SIH.git>



LANDING PAGE LINK

<https://SIH.lenlu.me>