

EDGE VIDEO ANALYTICS ACAP · AXIS CAMERAS

Device VDO

Detection, capture & recognition that run on the camera.

Device VDO turns any supported Axis camera into a self-contained analytics node. It detects people, faces and vehicles directly on the device, applies your zones and rules, captures frame-accurate evidence to the SD card, and posts events to any backend you choose — cloud or on-prem. No external server required to run; integrate it with your platform, or ours.



On-device AI

People · Faces · Vehicles

Draw-on-frame zones

SD-card evidence

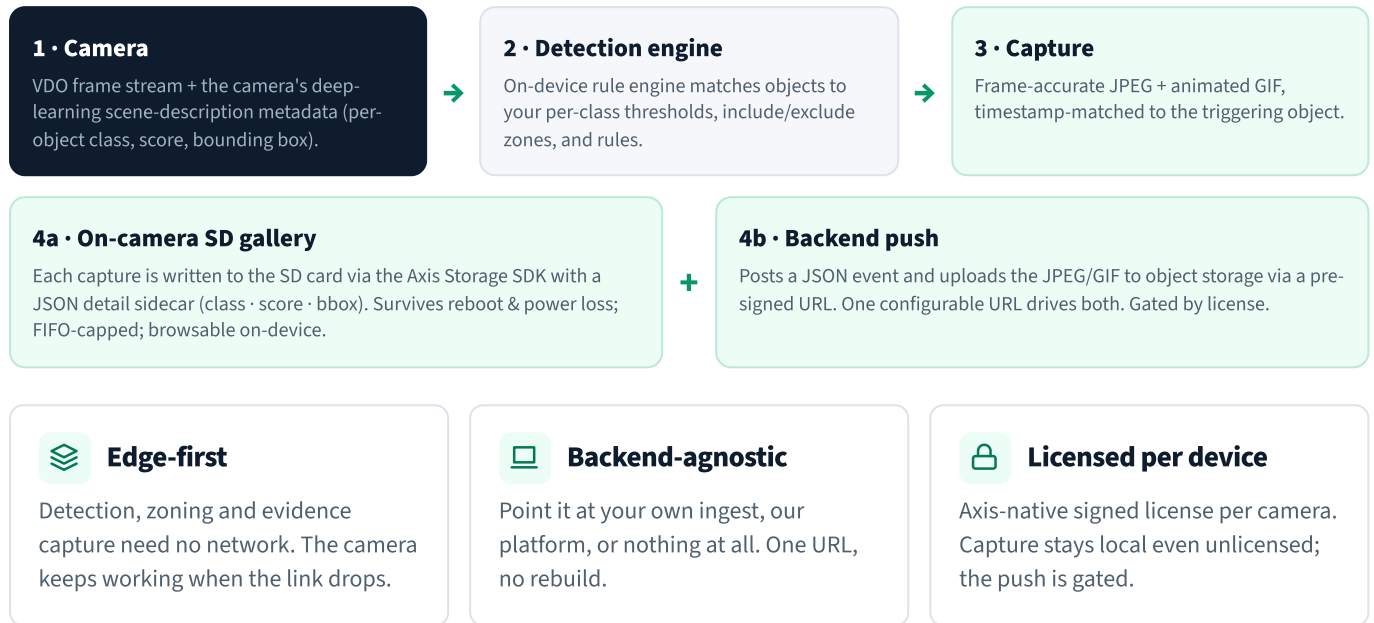
Backend-agnostic

Axis-native licensing

ARCHITECTURE

How the system works

Everything that decides *what* and *when* runs on the camera. The cloud only receives what you've decided is worth keeping — and only when a valid license is present.



Frame-accurate by design

Object metadata and video frames share the camera's monotonic clock. Device VDO matches each detection to the exact frame that produced it — the evidence shows the object that actually fired the rule, not a frame moments later.

DETECTION & RULES

What it can see — and decide

A configurable rule engine running on the camera, tuned from a built-in settings page. No code, no cloud round-trip.

Object classes

People, faces and vehicles (car / truck / bus / bike) from the camera's on-device deep-learning model. Each class has its own minimum-confidence threshold.

- Per-class sensitivity sliders (0–100%)
- Faces sourced on-device, recognized downstream

Detection zones

Draw include / exclude polygons directly on the live camera frame. An object counts only if its footprint is inside an include zone and outside every exclude zone.

- Per-zone class filter (e.g. people only)
- Normalized coordinates — resolution independent

Rules

Line crossing (with direction), time-in-area dwell, occupancy thresholds, and tailgating — each scoped to chosen classes with independent cooldowns.

- Directional line-crossing for in/out counting
- Per-rule confidence & size gates

Evidence capture

On every trigger: a full-resolution JPEG plus a short animated GIF of the surrounding moment, with the object's bounding box recorded for overlay.

- JPEG + multi-frame GIF clip
- Bounding box stored for review overlay

3

OBJECT CLASSES

4

RULE TYPES

2

ARCHITECTURES

0

SERVERS REQUIRED

STORAGE & INTEGRATION

Evidence on the edge, events anywhere



On-camera SD gallery

Snapshots are written to the SD card through the Axis Storage SDK — the real, mounted card, not volatile memory — so evidence survives reboots and power loss. A built-in web gallery shows each capture with its class, score and bounding-box overlay.

- FIFO-capped, automatic cleanup
- Per-snapshot JSON detail sidecar
- Works with zero backend configured



Backend integration

A single Events URL drives the whole outbound path: the device requests a pre-signed upload URL for the JPEG/GIF, uploads to object storage, then POSTs a JSON event. Repoint the URL to move the device between platforms — no rebuild.

- Per-request auth via device id + token
- Self-describing API index endpoint
- Standalone, or managed by Cube Access

The integration contract

Every call carries `?mac=&token=`; the consuming backend validates that pair. Endpoints derive from one base URL: `POST .../device/metadata` (event) and `POST .../createSignedUrl` (upload target). The on-camera settings page documents the live endpoint list.

OPTIONAL · CLOUD RECOGNITION

Familiar-face recognition

Captured faces flow to a recognition service that turns each face into a 512-dimension embedding (InsightFace / ArcFace) and matches it against your watchlist using vector similarity search. Matches raise a "familiar face" event with the person and a confidence score — offered as a managed cloud add-on, or run the embedding model yourself.



SPECIFICATIONS

Technical specifications

Platform	Axis ACAP — Native SDK 12.9 (compiled Go + C)
Architectures	aarch64 and armv7hf (one package per chipset)
Inputs	VDO video stream · on-device deep-learning scene-description metadata (objects, class, score, bounding box)
Detection	People · Faces · Vehicles — per-class confidence; include/exclude polygon zones; line-crossing, time-in-area, occupancy, tailgating rules
Capture	Full-resolution JPEG + animated GIF, frame-matched on the camera's monotonic clock
On-camera storage	SD card via Axis Storage SDK — reboot-persistent, FIFO-capped, JSON detail sidecar per snapshot, built-in gallery UI
Backend output	JSON event POST + pre-signed object-storage upload; single configurable Events URL; <code>?mac=&token=</code> auth
Configuration	On-camera settings page — sensitivity sliders, draw-on-frame zone editor, gallery, live API index
Recognition (optional)	ArcFace 512-d face embeddings + watchlist vector search (managed cloud add-on)
Licensing	Axis-native signed license — per device, per application; soft enforcement

Soft licensing

Install the signed license on the camera's Apps page. Without a valid license, on-device detection and the SD gallery keep running — only the backend push and configuration UI are gated. No data is lost; nothing breaks.

Privacy posture

Run fully standalone with evidence kept on the camera, push to your own infrastructure, or opt into managed cloud recognition. You choose where faces and events live.

Get Device VDO

Download the signed ACAP for your chipset at cube-access.com/downloads and install it from the camera's Apps page. Talk to us about licensing and managed recognition.