Analytics Tuning / False-Alarm Checklist - One-Page Reference

How to tune deployed video analytics so the alerts stay trustworthy: the vocabulary, the threshold dial, the false-alarm filter cascade, the base-rate trap, and the human-in-the-loop rules. Accuracy is a range you set, never a single number; aim for a stream an operator still trusts at hour seven, not for zero.

A. Accuracy is a range, not a number

- Every analytic makes two mistakes: a FALSE POSITIVE (false alarm - it fires when nothing happened) and a FALSE NEGATIVE (a miss - a real event ignored). You cannot drive both to zero, so 'accuracy' is a dial you set, not a number a vendor can promise. Reject any '100% accuracy' claim.
- Track two ratios, not one. PRECISION = TP / (TP + FP): of the alarms, how many were real. RECALL = TP / (TP + FN): of the real events, how many were caught. Decide which one your site cares about before you tune.

B. The threshold dial - pick the operating point

- The master dial is the CONFIDENCE THRESHOLD. Lower it for more recall (catch everything, more false alarms); raise it for more precision (clean alerts, more misses). The model is fixed - you are choosing a point on its curve.
- Set the point by the COST of each mistake. Perimeter/intrusion: a miss is catastrophic -> tune recall first, verify the noise. Retail counting / watch-lists: false positives corrupt data or harm a person -> tune precision. Do not ship the vendor default of 0.5 blindly.

C. Cut false alarms by stacking filters

- Do NOT just crank the threshold - that hides real events too. Stack filters, each removing one class of false alarm: object classification (person/vehicle only - the biggest cut, ~90% vs raw motion); detection zone (ROI); direction/line rule; object-size + perspective calibration; dwell/persistence (N seconds); arming schedule; rule combination with AND.
- Worked cascade: 40 cameras x 30 raw motion alarms = 1,200/day -> AI classification (-~90%) -> ~120/day -> zones + schedule + dwell (-~90%) -> ~12/day reviewable. Each filter targets noise, so recall survives.

D. The base-rate trap - rare events

- When the target event is RARE, a tiny error rate buries the true hits. 50,000 detections, 5 real events, 1% false-positive rate -> ~500 false vs ~5 true -> a '99% accurate' system is right ~1% of the time it alarms. The math cannot be argued with.
- So run rare-event and biometric analytics (watch-list face match, learned anomaly) as TRIAGE + VERIFICATION that cues a human - never as an autonomous alarm. Raise precision with a confirming second analytic, sensor, or operator.

E. The operator's reality - keep a human in the loop

- ALERT FATIGUE is the real failure mode: flood an operator with false alarms and they stop trusting and reading the system ('cry wolf'). Human attention on quiet video drops fast (~45% of activity missed after ~12 min, ~95% after ~22 min). Tune for trust, not zero.
- MEASURE to tune: track false alarms per camera per day and test recall with deliberate walk-tests; re-tune seasonally (foliage, lighting, traffic change). Keep a human confirming consequential alarms - GDPR Art. 22 (solely-automated decisions), EU AI Act Art. 14 (human oversight, automation bias), Art. 35 DPIA.

This is engineering guidance, not legal advice; confirm specifics with qualified counsel. The worked numbers are illustrative arithmetic (the 1,200 -> ~120 -> ~12 cascade and the 50,000 / 5 / 1% base-rate example depend on camera count, scene, and the real false-positive rate); the ~90% AI-vs-motion and ~90% verified-response figures are vendor/industry claims, not a benchmark; the vigilance figures are a widely cited security-industry figure with the phenomenon confirmed by the CCTV-vigilance literature. The detection-model internals are engineered in the AI for Video Engineering section; this is the surveillance APPLICATION - the operating point, the filters, the operator reality, the law. Sources: ONVIF Analytics Service Specification Ver. 25.12 + Profile M (rule config: CreateRules/ModifyRules, Annex A rule types); GDPR Reg. (EU) 2016/679 Art. 4(1)/22/35; EU AI Act Reg. (EU) 2024/1689 Art. 14; EDPB Guidelines 3/2019.