Behavioral Analytics (Loitering, Intrusion, Crowd, Zones) - One-Page Reference

What behavioral analytics is, the five rules, why false alarms are the real metric, how a behavioral event surfaces in the VMS, and why it is personal data but not biometric. Representative 2025-2026 figures - all move with scene, lighting, angle, and tuning.

A. What behavioral analytics is

- Behavioral analytics turns an object's POSITION and DWELL TIME into an alert. It answers geometry-and-timing questions - did something cross a line, enter a zone, stay too long, gather, or get left behind - not WHO someone is.
- It is not a separate AI: it is object detection + tracking (which build a 'scene description' of boxes, classes, and tracks) plus a geometry-and-timing RULE on top. Rule-based (you draw the line/zone) is the workhorse; learned behavioral-anomaly is a different tool (see anomaly detection).

B. The five rules you will actually use

- Line crossing (virtual tripwire): a tracked object crosses a drawn line, optionally in one direction. Intrusion zone (field): an object enters/stays inside a drawn polygon.
- Loitering: a zone rule with a dwell clock (commonly a few seconds to several minutes). Crowd / occupancy: count or density in an area, alarm on a threshold.
- Object left behind: an item that becomes and stays static (and its mirror, object removed) - the hardest of the five under occlusion and crowds.

C. False alarms are the real metric - and the math

- Behavioral systems watch quiet scenes where almost nothing is real, so the FALSE-ALARM rate, not the detection rate, decides success. Too many false alerts cause alarm fatigue - operators ignore or switch the analytics off.
- One outdoor camera on legacy pixel-motion: ~300 alerts/day, ~95% nuisance = ~285 false/day. Switch to a person-only, direction-aware tripwire: a ~90%+ cut -> ~29 false/day, most plausible.
- Well-tuned object-classified analytics target a false-positive rate in the low single digits, vs ~80%+ for raw motion. 'Well-tuned' is the load-bearing phrase: tuning is a deployment activity. The floor is never zero, and accuracy is never 100%.
- Biggest levers, in order: the object-CLASS filter (fire on people, not weather/animals), the GEOMETRY (place the line where only real crossings happen), the SCENE (rain/fog/low light), the angle/distance, and the thresholds.

D. How it surfaces - the ONVIF standard

- The ONVIF Analytics Service Specification defines the rule types as a normative rule engine: a Line Detector (-> Crossed event), a Field Detector (-> ObjectsInside), a Loitering Detector, and counting rules; a ClassFilter restricts a rule to object classes.
- A VMS can ask a camera which rules it supports, create/modify them, and subscribe to their events, which surface via ONVIF Profile M (over the stream, the event service, or MQTT) - so a behavioral alert from one vendor's camera works in another vendor's VMS.
- ONVIF standardizes the rule TYPE and the EVENT - not the detection model, the tracking, or the tuning that set the real false-alarm rate; those live in the vendor SDK. Conformance is a baseline for interoperability, not a guarantee of accuracy.

E. The legal line: personal data, but NOT biometric

- Behavioral analytics on identifiable people is PERSONAL DATA under GDPR (Art. 4(1)) - so you need a lawful basis, notice, and a DPIA for systematic monitoring of a public area (Art. 35; EDPB Guidelines 3/2019).
- But counting, zones, line-crossing, and loitering measure WHERE an anonymous object goes, not WHO it is - so they are NOT Art. 9 biometric data. The EDPB ranks simple counting as a less intrusive technology. Lighter than face recognition.
- It tips into the heavier regime the moment it IDENTIFIES a person or infers emotion/intent (face watchlist; EU AI Act Art. 5 restricts emotion recognition and real-time remote biometric ID). Keep behavior ANONYMOUS by design: count, don't identify.

Accuracy is a range tied to scene, lighting, angle, and tuning, never a single number and never 100%. Line and zone are the most reliable; loitering is all in the threshold; crowd counting is an ESTIMATE (dense-scene miscounts of dozens are normal - use it for thresholds, not exact headcounts); object-left-behind is hardest and needs measured proof (the UK i-LIDS benchmark certifies exactly these scenarios). The detection/tracking MODELS are engineered in the AI for Video Engineering section; this is the surveillance APPLICATION - the rules, the VMS authoring, the standard, and the law. Sources: ONVIF Analytics Service Specification (Annex A) + Profile M; GDPR Reg. (EU) 2016/679 Art. 4(1), 35; EDPB Guidelines 3/2019; EU AI Act Reg. (EU) 2024/1689 Art. 5; UK Home Office i-LIDS; crowd-counting MAE survey (ShanghaiTech, 2025); industry false-alarm analyses (representative, condition-dependent).