ONVIF Events & Metadata — Integration Cheat Sheet

Events say "something happened"; metadata describes "what is in the scene." This one-pager maps how each travels from a camera into a VMS, what Profile M guarantees, and the four traps that break analytics pipelines.

The two channels: what travels from camera to VMS

Channel	What it carries	How it reaches the VMS
Events	"Something happened" — motion, line cross, zone entry, a count change, a plate or face hit	ONVIF events service (pull or push); optionally MQTT/JSON
Metadata	"What is in the scene" — per frame: each object's box, center, class, and tracking id	VND.ONVIF.METADATA track over RTP, alongside the video

Events delivery: pull vs push (a device supports at least one)

Model	How it works	Use it when
Pull — PullPoint	VMS creates a subscription and calls PullMessages to fetch queued events	Across firewalls / NAT; the safe multi-vendor default
Push — Base Notification	VMS subscribes once; the camera sends a Notify the instant an event fires	You control the network and need lowest alert latency
Both	Ride OASIS WS-BaseNotification 1.3 topics; subscribe by topic to filter	Always — subscribe only to the topics you need

Profile M: mandatory vs optional — read the Declaration of Conformance

Feature	Status	Note
Metadata streaming (RTP)	**Mandatory**	The scene-description stream itself
Analytics service	**Mandatory**	Discover and configure analytics modules
MQTT / JSON events	**Optional**	The IoT bridge — verify it on the DoC
Rule engine	**Optional**	Line-cross, zone, counting rule authoring
Geolocation / face / plate	**Optional**	Rich attributes vary by firmware build

The four traps that break analytics pipelines: (1) archive the metadata — most VMS discard it, so forensic search is empty a week later; (2) sync clocks — NTP drift misaligns metadata with the video frame; (3) map class taxonomies — one vendor's "Human" is another's "Person"; (4) tame event storms — raise confidence thresholds, aggregate by object id, and filter by topic. ONVIF standardizes the format and transport, never the detection accuracy or the rule-authoring format. Sources: ONVIF Profile M Specification v1.1; ONVIF Analytics / Core / Streaming Specifications; OASIS WS-BaseNotification 1.3.