Privacy-Native Edge Video Analytics for Regulated Rail Environments

A Dhi Technologies whitepaper — Dev Sanghvi, Dhi Technologies · July 2026 (draft)

Executive summary

Rail operators are being asked to do two things at once: expand camera-based safety and compliance programs, and comply with a rapidly hardening data-protection regime — India's Digital Personal Data Protection Act (DPDP) with its Rules notified in November 2025, and in Europe the AI Act's staged obligations alongside GDPR's long-standing video-surveillance guidance. Most video-AI offerings answer this tension with paperwork: cloud platforms plus contractual assurances.

Dhi answers it with architecture. In the Dhi platform, **footage never leaves the premises**. Every analytic — object detection, safety-compliance checks, license-plate reading, and even vision-language scene understanding — runs on an edge device installed on-site. What crosses the network boundary is derived metadata: events, alerts, and anonymous identity tokens. Live video, when an authorized operator needs it, travels point-to-point through an authenticated, time-limited channel and is never stored or re-served from the cloud. Person continuity is built **face-free**: appearance-based tokens rather than facial identity, with any face workflow requiring explicit, per-subject enrollment.

The result is a system where the answers to a regulator's hardest questions — *where is the footage? who can see it? how long is it kept? is anyone biometrically identified?* — are properties of the deployment, not policies of a vendor. This paper explains the regulatory landscape as of mid-2026, the architecture, and an obligation-by-obligation mapping between the two.

---

1 The regulatory moment

1.1 India: DPDP is now operational

The Digital Personal Data Protection Act, 2023 became fully operational when the DPDP Rules were notified in November 2025, opening a phased compliance window of roughly eighteen months for organizations [1][2]. Three points matter most for video programs:

- **CCTV is permitted — with obligations attached.** The Act's "legitimate uses" (Section 7) allow processing without fresh consent for State instrumentalities delivering services under law (7(b)) and **for purposes of employment or safeguarding the employer from loss or liability (7(i))** — the ground under which workplace safety and hygiene monitoring typically proceeds [3]. The question DPDP asks of a rail operator is therefore not *whether* cameras may run, but *how* the resulting personal data is minimized, secured, retained, and erased.
- **The duties are architectural.** Data Fiduciaries must implement reasonable security safeguards (Section 8(5)), notify the Data Protection Board and affected individuals of breaches (8(6)), and erase personal data once its purpose is served (8(7)) [4].
- **The penalties are material.** The Act's schedule provides penalties up to ₹250 crore for failing to implement reasonable security safeguards and up to ₹200 crore for failing to notify a breach [1][4].

A centralized cloud archive of raw station or kitchen footage is, in DPDP terms, a large standing liability: a single breach surface aggregating personal data whose purpose was served seconds after analysis. An architecture that never aggregates footage shrinks the Section 8 exposure to metadata.

1.2 Europe: the AI Act draws lines around biometrics

Regulation (EU) 2024/1689 (the AI Act) is in staged application: its prohibitions have applied since 2 February 2025, and — following the 2026 "Digital Omnibus" simplification package — the compliance deadline for standalone high-risk systems under Annex III moved from August 2026 to **2 December 2027** [5][6][7].

For video analytics the classification lines are:

- **Prohibited:** real-time remote biometric identification in publicly accessible spaces for law enforcement, with narrow, authorization-gated exceptions (Article 5(1)(h)) [5].
- **High-risk (Annex III):** remote biometric identification, biometric categorisation on sensitive attributes, and emotion recognition; separately, systems for critical infrastructure safety, and — often overlooked — **systems that monitor and evaluate the performance and behaviour of workers (Annex III 4(b))** [6].
- **Outside those categories:** analytics that identify *situations rather than people* — PPE presence, intrusion into a hazardous zone, fire and smoke, hygiene-protocol compliance — are not biometric systems on their face.

Two honest nuances belong in any credible deployment plan. First, workplace analytics (for example, kitchen hygiene compliance) may be argued to fall under Annex III 4(b) as worker monitoring even when fully non-biometric; classification deserves case-by-case legal review, and an operator should assume high-risk obligations (risk management, data governance, logging, human oversight) *might* attach [6]. Second, biometric *verification* (1:1) is explicitly excluded from the high-risk biometrics category — relevant where staff authentication is contemplated [6]. In both cases the practical posture is the same: an architecture that already minimizes data, logs decisions, keeps humans in the loop, and avoids biometric identification by default makes whichever classification lands survivable.

1.3 GDPR's video-surveillance doctrine

The European Data Protection Board's Guidelines 3/2019 on video devices remain the reference for CCTV data protection: purposes must be specified before deployment (vague "safety" is insufficient), retention should be days not months absent specific justification, transparency requires layered notices, and facial recognition through CCTV triggers heightened requirements [8]. One doctrinal point does a lot of work: **ordinary video only becomes special-category biometric data when processed through specific technical means for uniquely identifying a person** (GDPR Articles 4(14) and 9(1)) [9]. A system that never runs facial identification on its streams keeps its video processing outside Article 9's heightened regime — this is the legal footing of the face-free design in §3.4.

1.4 Rail context

Indian Railways has been expanding station video surveillance for years — an IP-based video surveillance program spanning nearly a thousand stations under the Nirbhaya Fund [10] — and

railway catering operations run continuous CCTV monitoring of food preparation alongside food-safety certification requirements [11]. The direction of travel is more cameras and more analytics inside an increasingly regulated data regime: exactly the squeeze this architecture is designed for.

---

2 Why cloud-first video AI struggles here

A conventional VSaaS pipeline uploads camera streams to cloud storage and runs analysis there. Under the regime above, that design imports four structural problems:

1. **It maximizes the personal-data footprint.** Footage whose analytical purpose is served in seconds persists in a remote archive, expanding breach exposure (DPDP §8(5)/(6); GDPR storage limitation) in proportion to fleet size and retention length.
2. **It centralizes the breach surface.** One compromised tenant boundary exposes many sites' footage; penalties now attach to precisely this failure.
3. **It complicates purpose limitation.** Archived raw video invites secondary uses that were never specified at collection — the exact pattern the EDPB guidance targets.
4. **It carries a bandwidth and cost tax.** Continuous multi-camera upload consumes uplink capacity and recurring storage spend that industry estimates place in the tens of thousands of dollars per month for large fleets — figures that vary by source but point one direction [12].

None of this makes cloud video AI unlawful. It makes it *expensive to defend*. The alternative is to not create the liability.

---

3 The privacy-native architecture

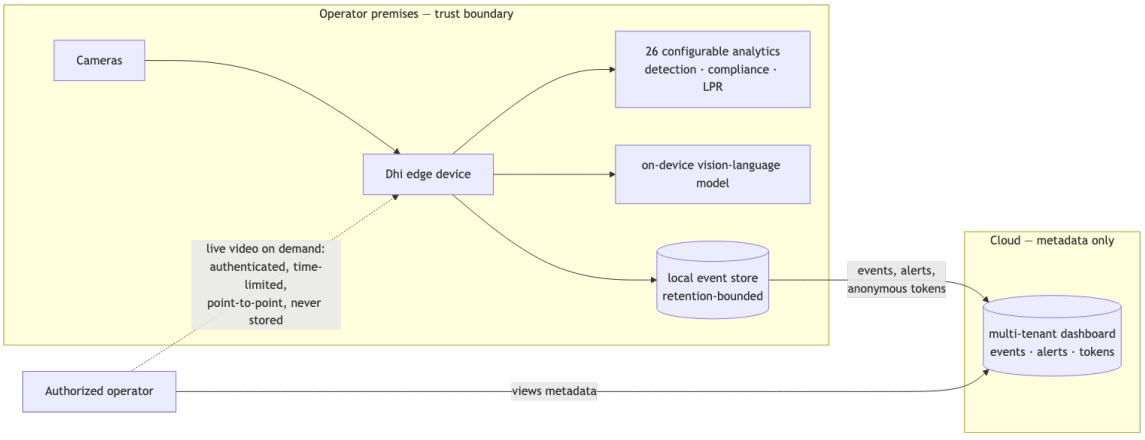

Figure 1: Figure 1. What crosses the boundary is the design.

Figure 1. What crosses the boundary is the design. Footage and frames stay on-premises; only derived metadata goes up; live video comes down point-to-point on demand.

3.1 All inference on the premises

The Dhi edge device — a small NVIDIA Jetson-class unit — runs the full analytics stack for its cameras: a shared-inference architecture (documented separately in our systems paper) that fits detection, OCR, and policy logic for six concurrent streams in a fraction of an 8 GB device, with 26 canonical use cases across 11 detector engine families enabled per site by configuration. Frames are analyzed in memory on the device; they are not shipped elsewhere for inference.

3.2 Even the "AI questions" stay on-box

Where scene understanding beyond detection is wanted ("what is happening near the loading door?"), the device runs a quantized multimodal vision-language model locally. No frame is sent to a third-party model API. This closes the most commonly overlooked leak in modern video stacks: the analytics may be on-prem while the "smart" layer quietly posts screenshots to an external service. Here the answer to "which processors see the footage?" is: the device, full stop.

3.3 Metadata up, media down

The cloud dashboard receives *events* — "PPE violation, camera 3, 14:02, snapshot reference, confidence band" — plus alerts, device health, and anonymous continuity tokens. Live video, when an authorized operator requests it, streams from the device through an authenticated tunnel with time-window-signed access, point-to-point to the viewer's browser; the cloud never stores or re-serves it. If the device is offline, live video is simply unavailable — a deliberate trade of convenience for the guarantee that no cloud copy exists.

3.4 Face-free by default

Person continuity across cameras uses appearance-based tokens rather than facial identity, and vehicle continuity uses normalized license plates — identifiers scoped to the operational question ("did the same person/vehicle appear at camera 1 and camera 4?") rather than to *who someone is*. Any face-based workflow requires explicit, per-subject enrollment; there is no ambient face database and no real-time remote biometric identification, keeping the system clear of the AI Act's prohibited category and, under GDPR doctrine, outside special-category processing for ordinary streams [6][9]. We present face-free operation as an architectural commitment built on the mature appearance-based re-identification literature [15], not as a single off-the-shelf technique.

3.5 Governance as code

Retention windows are enforced on the device (events age out; storage is bounded). Access is organization-scoped end-to-end: an operator authenticates through a zero-trust identity layer, and every query is filtered to their organization's rows. Ingestion requires authenticated internal channels. Audit trails record who saw what. These are defaults in the platform, not integration projects.

---

4 Mapping obligations to architecture

Obligation (regime)	Architectural answer
Data minimization (DPDP §4/purpose; GDPR Art. 5; EDPB 3/2019)	Only derived metadata leaves the site; frames analyzed in memory; snapshots stored locally, referenced remotely
Storage limitation (GDPR; EDPB "days, not months"; DPDP §8(7) erasure)	Device-enforced retention windows on events and media; no cloud footage archive to govern
Security safeguards (DPDP §8(5))	No centralized footage honeypot; org-scoped access; authenticated ingestion; per-device credentials
Breach notification exposure (DPDP §8(6))	Worst-case cloud breach discloses event metadata, not video; footage breach requires physical/site compromise
Purpose limitation (EDPB 3/2019)	Analytics are explicit per-site configuration (use-case registry); enabling a new purpose is a logged, auditable act
No prohibited biometric ID (AI Act Art. 5(1)(h))	No real-time remote biometric identification capability in the deployed configuration; face-free continuity by default
High-risk readiness if Annex III attaches (e.g., 4(b) worker monitoring)	Logging, human-in-the-loop alert review, documented data governance, bounded retention — the high-risk obligations' substance is already present
Transparency (EDPB layered notices; DPDP notice)	Operator-facing signage/notice templates; per-use-case documentation of what is detected and what is stored
Data Principal rights (DPDP §11–13: access, correction, erasure)	Metadata is queryable and erasable per subject token; no sprawling raw-footage estate to search

Table 1. The claim is not "compliance by default" — deployments still need notices, DPIAs where applicable, and legal review. The claim is that the architecture makes each obligation cheap to meet and easy to evidence.

5 Deployment scenarios

Food-safety compliance in railway base kitchens. A national railway catering operator's base kitchens require continuous verification of hygiene protocols: PPE and attire compliance, surface and process checks, presence in prohibited zones, contamination-risk events. The Dhi configuration runs these as on-device analytics with alert-tiered escalation; supervisors see events and annotated snapshots, not a raw feed warehouse; workers' footage never leaves the kitchen. Because this is workplace monitoring, the deployment posture assumes worker-management obligations may apply: alerts are human-reviewed, detection scope is documented per camera, and retention is short and enforced. **[Data from the first kitchen pilot will be reported separately when available.]**

Station and passenger-safety analytics. Intrusion on track aprons, crowd density, abandoned objects, fire/smoke — situation analytics that require no identity at all. The face-free default means expanding camera counts does not expand biometric risk.

Multi-site fleets. Each site's device is claimed into the operator's tenant via a short-lived pairing code; per-device credentials are hashed at rest; a site can be onboarded in minutes without exposing any site's data to any other tenant.

---

6 What we deliberately do not do

- No cloud storage or re-serving of footage — the cloud tier cannot leak what it does not hold.
- No ambient facial recognition and no face database without explicit enrollment.
- No third-party video or VLM APIs in the inference path.
- No silent repurposing: analytics are a declared, per-site configuration surface.

---

7 Roadmap

The same architecture extends to: uncertainty-aware alerting (verifying borderline alerts with the on-device VLM before an operator sees them); deeper face-free re-identification as the appearance-based literature matures [15]; and federated improvement loops where models improve from on-site hard cases without raw footage ever pooling centrally — a direction consistent with recent federated video-analytics research [13][14].

References

[1] Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023), Government of India; DPDP Rules notified November 2025 (Gazette G.S.R. 846(E)). See MeitY/PIB notifications. [2] Press Information Bureau, "DPDP Rules 2025 notified," November 2025. [3] DPDP Act §7 ("certain legitimate uses"), incl. 7(b) State functions, 7(i) employment. [4] DPDP Act §8 (Data Fiduciary obligations: 8(5) safeguards, 8(6) breach notification, 8(7) erasure) and Schedule (penalties). [5] Regulation (EU) 2024/1689 (AI Act), Article 5(1) prohibited practices; prohibitions applicable 2 Feb 2025. artificialintelligenceact.eu/article/5/. [6] AI Act Annex III (high-risk categories, incl. 1 biometrics, 4(b) worker monitoring; biometric-verification exclusion). artificialintelligenceact.eu/annex/3/. [7] Council of the EU / European Parliament, Digital Omnibus on AI: high-risk (Annex III) compliance deadline deferred to 2 Dec 2027 (Parliament endorsement 16 Jun 2026; Council approval 29 Jun 2026). [8] EDPB, Guidelines 3/2019 on processing of personal data through video devices. [9] GDPR Art. 4(14), Art. 9(1) — biometric data and special-category processing. [10] Reporting on Indian Railways' IP-based video surveillance program (983 stations, Nirbhaya Fund); program scope figure, completion status varies by report date. [11] Trade reporting on railway base-kitchen CCTV and FSSAI-linked food-safety practice. [Primary ministry circular citation pending — flagged, not asserted.] [12] Industry market estimates on surveillance storage/bandwidth economics (directional; sources vary materially). [13] Cangialosi et al., "Privid: Practical, Privacy-Preserving Video Analytics Queries," USENIX NSDI 2022, arXiv:2106.12083. [14] Poddar et al., "Visor: Privacy-Preserving

Video Analytics as a Cloud Service," USENIX Security 2020, arXiv:2006.09628; plus 2025 federated video-analytics work (e.g., arXiv:2510.17651, arXiv:2511.07171). [15] Appearance-based / cloth-changing person re-identification literature, e.g. Pang et al., "Identity-Clothing Similarity Modeling for Unsupervised Clothing Change Person Re-Identification," CVPR 2025; survey context in "Vision-Language Models for Edge Networks," arXiv:2502.07855.

Contact: Dhi Technologies · dhi-tech.com