Network Working Group
Internet-Draft
Intended status: Informational
Expires: 30 November 2026

L. Melegassi
Catellix

29 May 2026

MVPS Video-Surveillance and CCTV Profile: Fleet-Coherence Detection of Feed Replay and Loop Injection Across IP Cameras

draft-melegassi-ippm-mvps-video-surveillance-00

Abstract

This document defines a Multi-Vantage Path Snapshot (MVPS) domain profile for video surveillance: fleets of IP cameras, NVR/VMS recorders, and cloud video-surveillance-as-a-service (VSaaS) endpoints treated as MVPS vantages. Its target threat is the feed-replay (loop) attack -- an on-path adversary that substitutes a live camera stream with previously recorded footage so that operators and single-camera analytics see nothing wrong.

The profile re-establishes the bounded-joint-skew axiom A1 under the video pipeline (camera NTP/PTP residual, encoder and GOP buffering, recorder jitter buffer, frame-interval quantization), gives a closed-form maximum pipeline budget, and proves the headline result: with capture timestamps authenticated at the sensor, any replayed loop older than a closed-form, strictly positive minimum age leaves the coherence ball and is flagged by core Theorem T2. The core detection and Byzantine theorems are inherited via the MVPS Architecture-Invariance Theorem.

The profile is DEFENSIVE: it detects coherence anomalies (feed replay, loop, tamper, rogue ingest). It defines no facial recognition, biometric, tracking, or identification function. All properties are validated by scripts/validate_video_surveillance.py (7/7 PASS, exit 0) and recorded in evidence/video_surveillance_receipt.json.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 30 November 2026.

Table of Contents

1. Introduction

MVPS detects network-propagating anomalies by measuring the coherence of an observed state across multiple spatially independent vantages. Its theorems are surface-independent: they hold where the five MVPS axioms hold, by the Architecture-Invariance Theorem [I-D.melegassi-iab-mvps-architecture].

A video-surveillance deployment is such a surface. Every IP camera, NVR/VMS recorder, and cloud-VSaaS endpoint is a vantage that emits timestamped state on a tick. Camera fleets are large, long-lived, physically distributed, and a frequent target -- which makes them a natural and high-value instantiation of MVPS.

1.1. The Feed-Replay Threat

The signature attack against surveillance is the feed-replay (loop): an on-path adversary substitutes the live stream with previously recorded footage so an operator -- or a per-camera analytic -- sees a plausible but stale scene while something happens off-frame.

Classical countermeasures are per-camera (motion entropy, frame hashing, watermarks). They are blind to the one asset a fleet has that a single camera does not: COHERENCE with the rest of the fleet's clocks. A replayed feed is, by construction, late: its authenticated capture time lags real time by the loop's age. MVPS measures exactly that lag as a coherence offset across the fleet.

1.2. Defensive Scope and Non-Goals

This profile is strictly DEFENSIVE: detection of coherence anomalies in camera-fleet telemetry (feed replay, loop, tamper, rogue ingest).

This document does NOT define and MUST NOT be claimed to define any facial recognition, biometric, person-tracking, or identification function, nor any output other than coherence-anomaly detection and audit logs.

2. Terminology

eps_ntp: the camera capture-timestamp residual (NTP/PTP).

tau_enc: encoder plus GOP/keyframe buffering latency.

tau_jb: jitter buffer at the VMS/NVR/cloud ingest.

tau_frame: frame-interval quantization, = 1/fps.

delta_replay: the age of an injected loop / replayed segment.

T_tick: the deployment coherence tick.

Authenticated capture time: a capture timestamp bound at the sensor (signed frame or attested encoder) so it cannot be silently rewritten on the wire.

The key words "MUST", "MUST NOT", "SHOULD", "MAY" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals.

3. The Video-Pipeline Joint-Skew Model

A disciplined enterprise camera holds capture time to eps_ntp via PTP (about 1e-4 s); a cheap IP camera over NTP holds 1e-2..1e-1 s. Encoding with a GOP/keyframe structure buffers up to tau_enc before a coherent frame is emitted; the recorder or cloud ingest adds a jitter buffer tau_jb. Sampling at f frames per second quantizes capture time by tau_frame = 1/f. The effective joint skew is

$$skew_{eff} = 2*eps_{ntp} + tau_{enc} + tau_{jb} + tau_{frame} .$$

The replay attack is treated in Section 6 as an additive offset to this skew.

4. Re-establishing Axiom A1 (Lemma L-CAM-1)

Axiom A1 holds on tick T_tick iff

$$skew_{eff} = 2*eps_{ntp} + tau_{enc} + tau_{jb} + tau_{frame} < T_{tick}.$$

For representative budgets:

- enterprise-PTP (ONVIF Profile-T, low-latency H.265):
  skew_eff = 63.5 ms < 1000 ms tick
- consumer-IP camera (NTP, GOP buffering):
  skew_eff = 466.7 ms < 1000 ms tick
- cloud-VSaaS (server transcode + ingest buffer):
  skew_eff = 880.0 ms < 2000 ms tick

All satisfy A1; the infeasible control (2.5 s of buffering against a 1 s tick) gives skew_eff = 2667 ms and correctly violates A1 (validator check L-CAM-1).

5. Maximum Pipeline Budget (Lemma L-CAM-2)

Solving skew_eff = T_tick for the end-to-end buffering gives

$$\tau_{pipe,max} = T_{tick} - \tau_{frame} - 2\,\varepsilon_{ntp}.$$

For the enterprise-PTP budget, $\tau_{pipe,max}$ is about 966.47 ms at a 1 s tick. The practical reading is that the binding term is the pipeline buffering (encoder GOP plus recorder jitter buffer); keeping it under $\tau_{pipe,max}$ preserves coherence at the chosen tick.

6. Feed-Replay Detectability (Lemma L-CAM-4)

Suppose capture timestamps are authenticated at the sensor, so the adversary cannot rewrite them without detection by the authentication layer. An injected loop of age delta_replay raises the replaying vantage's apparent skew to skew_eff + delta_replay. The replaying vantage therefore leaves the coherence ball -- and is flagged by core Theorem T2 -- as soon as

$$\Delta_{replay} \geq \Delta_{min} = T_{tick} - \text{skew}_{eff} \quad \text{(the A1 margin)}.$$

This minimum detectable loop age is closed-form and strictly positive on every feasible budget:

enterprise-PTP	Δ_{min} = 936 ms	(at a 1 s tick)
consumer-IP	Δ_{min} = 533 ms	(at a 1 s tick)
cloud-VSaaS	Δ_{min} = 1120 ms	(at a 2 s tick)

A tighter tick lowers Δ_{min} and catches shorter loops, trading detector sensitivity against pipeline headroom (validator check L-CAM-4: each Δ_{min} is exact -- skew_eff + Δ_{min} round-trips to T_{tick}).

7. Inheritance of the Core Theorems

If A1 holds (Section 4) and the compromised-vantage fraction $f < 1/2$, then by the Architecture-Invariance Theorem [I-D.melegassi-iab-mvps-architecture] the core results inherit verbatim:

T1	multi-vantage D^2 dominates per-vantage max-z;
T2	Φ_D concentration under the null;
T3′	empirical-quantile false-alarm calibration;
T9	Byzantine robustness of the geometric-median aggregator.

No core theorem is re-derived (validator check A-CAM-INHERIT).

8. Byzantine and Compromised Cameras

A large camera fleet must assume some vantages are compromised (hijacked cameras, rogue ingest, botnet-recruited DVRs). For $f < 1/2$ the geometric-median aggregator has finite max-bias $b(f) = C f/(1-2f)$ (after [Minsker]; MVPS imported result I12), diverging only as $f \to 1/2$ (validator check B-CAM-1: $b(0.2)=0.333$, $b(0.4)=2.000$).

9. Timestamp-Blind Loops (Conjecture C-CAM-1)

When capture timestamps are NOT authenticated (the adversary rewrites them along with the loop), it is plausible that a loop is still

detected via cross-vantage SCENE coherence -- overlapping fields of view, shared illumination / shadow dynamics, correlated motion phase -- flagged by the multi-vantage detector before any single-camera tamper analytic alarms. This is stated as a CONJECTURE, not a theorem, with a falsification protocol (observable: cross-vantage correlated luminance / motion-phase anomaly vs per-camera tamper score; data: an overlapping-FOV fleet with ground-truth injected loops; test: Wilson 95% lower bound on detection-time gain > 0; blocker: a calibrated overlapping-FOV testbed with labelled replay injections). The headline replay defence (Section 6) stands on its own ONLY with authenticated timestamps; this conjecture is its unproved timestamp-blind complement, and the profile's guarantees do NOT depend on it.

10. Operational Logging

Deployments SHOULD log events using the MVPS operational log format [I-D.melegassi-opsawg-mvps-logging], anchoring opportunistically; a flagged replay/loop event and the offending vantage's offset are themselves useful, tamper-evident audit records.

11. Numerical Receipt

scripts/validate_video_surveillance.py evaluates seven checks (L-CAM-1..4, A-CAM-INHERIT, B-CAM-1, C-CAM-1) over the budgets above and writes evidence/video_surveillance_receipt.json with per-scenario skew, the closed-form pipeline tolerance, the per-scenario minimum detectable loop age, the inherited theorem list, the defensive non-claims, and a SHA-256 of its own canonical body. All seven checks PASS (exit 0).

12. Security Considerations

The profile is a detection and audit capability; no facial recognition, biometric, or targeting surface is added. Its value is early, coherent detection of feed replay, loop injection, tamper, and rogue ingest across a camera fleet, with a tamper-evident audit trail (Section 10).

The headline replay defence (Section 6) depends on authenticated capture timestamps; without them, loop detection is only a conjecture (Section 9) and MUST NOT be relied upon as a guarantee. Quantum-era integrity of logs and anchors follows the Proof Envelope [I-D.melegassi-ippm-mvps-proof-envelope].

13. IANA Considerations

This document has no IANA actions.

14. References

14.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017.

[I-D.melegassi-iab-mvps-architecture] Melegassi, L., "MVPS Architecture Invariance",

draft-melegassi-iab-mvps-architecture-00, 2026.

14.2. Informative References

[I-D.melegassi-ippm-mvps-orbital-coherence]
Melegassi, L., "MVPS Orbital Coherence", draft-melegassi-ippm-mvps-orbital-coherence-00, 2026.

[I-D.melegassi-ippm-mvps-maritime-edge]
Melegassi, L., "MVPS Maritime and Tactical-Edge Profile", draft-melegassi-ippm-mvps-maritime-edge-00, 2026.

[I-D.melegassi-ippm-mvps-terrestrial-mobile]
Melegassi, L., "MVPS Terrestrial Mobile and Vehicular Profile", draft-melegassi-ippm-mvps-terrestrial-mobile-00, 2026.

[I-D.melegassi-opsawg-mvps-logging]
Melegassi, L., "The MVPS Operational Log Format", draft-melegassi-opsawg-mvps-logging-00, 2026.

[I-D.melegassi-ippm-mvps-proof-envelope]
Melegassi, L., "MVPS Proof Envelope", draft-melegassi-ippm-mvps-proof-envelope-00, 2026.

[Minsker] Minsker, S., "Geometric median and robust estimation in Banach spaces", Bernoulli 21(4), 2015.

Appendix A. Worked Budgets (Normative)

The three budgets of Section 4 (enterprise-PTP, consumer-IP, cloud-VSaaS) and the infeasible control (2.5 s of buffering at a 1 s tick) are the normative vectors. A conformant implementation MUST reproduce, for each, the skew_eff value, the A1 verdict, and the minimum detectable loop age delta_min emitted by scripts/validate_video_surveillance.py.

Author’s Address

Leonardo Melegassi
Catellix
Brazil
Email: melegassi@catellix.com