INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-04-04 13:38 UTC

Unverified Claim: CISA Emergency Directive 26-34 on AI Video Surveillance 'Pixel Blindness' Attack

SECURITY ANALYSIS | LOW

SCC Item ID	SCC-STY-2026-0047
Type	Security Analysis
Severity	LOW
Affected Products	AI-driven video surveillance systems (unverified scope)
Published	2026-04-03
Discovery Source	Gemini

Executive Summary

A circulating claim describes CISA Emergency Directive 26-34 mandating audits of AI-driven video surveillance systems in response to an adversarial technique called 'pixel blindness', but this directive does not appear in CISA's published Emergency Directive registry, and the named attack technique has no verified presence in adversarial ML literature. This item cannot be confirmed as real and should not be treated as an actionable regulatory obligation. The underlying threat concept, adversarial manipulation of computer vision models to create physical detection blind spots, is legitimate and documented in adversarial ML research, making this a signal worth tracking even if the specific claim is fabricated.

Technical Analysis

INTEGRITY FLAG, UNVERIFIED/LIKELY FABRICATED CLAIM

As of the configuration date of 2026-03-04, no CISA Emergency Directive numbered 26-34 appears in CISA's publicly maintained Emergency Directive registry at cisa.gov/emergency-directives. CISA Emergency Directives are legally binding orders issued under 44 U.S.C. § 3553(h) and are published publicly upon issuance. A directive of this type mandating AI surveillance audits would be a significant regulatory action with a traceable public record. None exists here.

The named attack technique, 'pixel blindness,' does not correspond to any established, peer-reviewed, or industry-documented adversarial ML attack category known in the literature as of this configuration date. The five sources attached to this item are general AI surveillance industry marketing content with no connection to the claimed directive or technique.

What is verifiably true and analytically relevant: adversarial attacks against computer vision and AI-based detection systems are a documented and active research area. MITRE ATLAS catalogs adversarial ML tactics including evasion techniques targeting image classifiers and object detection models. Researchers have demonstrated that physical-world adversarial inputs, such as specially patterned clothing, IR light sources, or manipulated environmental features, can cause AI detection models to fail to classify or locate a subject. These techniques exploit the statistical nature of neural network inference rather than traditional software vulnerabilities. The concept of exploiting spatial reasoning dead zones in AI surveillance models is plausible as a research hypothesis, but no published, named technique called 'pixel blindness' is documented here.

The discovery source for this item was a broad-coverage AI search tool, not an authoritative regulatory or research source. This origin, combined with the absence of any verifiable directive, named technique, or supporting sourcing, is consistent with AI hallucination of plausible-sounding regulatory content.

Security teams should not action this as a compliance requirement. Organizations that rely on AI-driven video surveillance for physical security should independently assess adversarial ML exposure, not because of this directive, which cannot be confirmed, but because the underlying risk class is real. The appropriate reference framework is MITRE ATLAS, not the claimed directive.

Action Checklist

1. Step 1: Verify independently, check cisa.gov/emergency-directives directly to confirm whether any AI surveillance directive exists before treating this as a compliance obligation
2. Step 2: Do not redistribute, avoid forwarding this claim internally or to leadership as confirmed regulatory news; it has not been verified and may be fabricated
3. Step 3: Assess AI surveillance exposure on its merits, inventory AI-driven video analytics systems in your physical security stack regardless of this claim, as adversarial ML risk to computer vision is a documented threat class
4. Step 4: Reference MITRE ATLAS, review adversarial ML evasion techniques relevant to image classification and object detection systems using the ATLAS framework at atlas.mitre.org
5. Step 5: Monitor CISA channels, subscribe to CISA's official directive and alert feeds to receive confirmed regulatory actions at the source, reducing exposure to fabricated or misattributed claims

IR / Forensic Enrichment

Triage Priority	DEFERRED
Escalation Criteria	Escalate to leadership and legal counsel only if CISA independently publishes a verified Emergency Directive addressing AI video surveillance systems, if internal distribution of the unverified claim has already reached executive stakeholders or external parties creating potential regulatory misrepresentation risk, or if evidence emerges suggesting the fabricated directive is part of a coordinated influence operation targeting your sector's compliance posture.

Recovery Notes	Recovery for this item is primarily procedural rather than technical: update the threat intelligence intake process to require source verification against cisa.gov/emergency-directives before any regulatory claim enters internal distribution channels. Monitor CISA's official feeds for 30 days for any legitimate directive addressing AI surveillance systems that may have been anticipated by whoever fabricated this claim. Review whether any premature compliance actions (vendor calls, policy drafts, budget discussions) were initiated based on the unverified claim and document their reversal to maintain audit trail integrity.
Forensic Artifacts	Original claim artifact with full provenance: screenshot or PDF of the source (social media post, email, forwarded message, news article) including URL, publication timestamp, author or account name, and engagement metrics if applicable — establishes the disinformation vector and spread velocity | Internal distribution chain record: email thread headers (From, To, CC, Received chain), Slack or Teams message metadata, or ticket audit trail showing who received the claim internally, at what time, and what actions they initiated — required to assess organizational exposure and reverse any premature compliance actions | CISA Emergency Directive registry screenshot: timestamped capture of https://www.cisa.gov/emergency-directives at the time of verification showing absence of 'ED 26-34' or any AI surveillance directive — constitutes the primary evidence of fabrication and protects the organization if later questioned about regulatory due diligence | AI video analytics asset inventory output: nmap scan results, DHCP lease exports, and ARP table captures from the physical security VLAN documenting all camera endpoints, NVR hosts, and AI inference nodes discovered during Step 3 exposure assessment — establishes the actual attack surface for adversarial ML threats independent of the fabricated directive | Threat intelligence source registry update record: documentation of new CISA subscription confirmations (GovDelivery confirmation email, RSS feed URL) and any updates to the organization's threat intelligence intake procedure — demonstrates procedural improvement and supports NIST IR-8 plan update requirements

Per-Action IR Details

Step 1: Verify independently — check cisa.gov/emergency-directives directly to confirm whether any AI surveillance directive exists before treating this as a compliance obligation

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: validating whether an event constitutes a real incident before committing response resources

Controls: NIST IR-4 (Incident Handling) — requires verifying incident characteristics before initiating formal response actions, NIST IR-6 (Incident Reporting) — prohibits escalating unverified events as confirmed incidents to leadership or external parties, NIST SI-5 (Security Alerts, Advisories, and Directives) — mandates receiving directives from authoritative external organizations directly, not through secondary or unverified channels, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — verification of threat claims is prerequisite to initiating any compliance-driven remediation workflow

Compensating: Navigate directly to https://www.cisa.gov/emergency-directives in a browser — the page lists all active Emergency Directives chronologically with official directive numbers and publication dates. Cross-reference against CISA's @CISAgov official social channels and the Federal Register for any corresponding rulemaking. This takes under 5 minutes and requires no tooling. If 'ED 26-34' or any AI surveillance directive is absent from that page, treat the claim as unverified. Document your check with a screenshot timestamped via your OS clock.

Evidence: Before concluding verification, preserve the original claim artifact: screenshot or PDF the source where this directive claim appeared (social media post, email, Slack message, news article), capture the URL with a timestamp, and note the distribution chain (who forwarded it and from where). This establishes provenance for any future disinformation tracking or vendor/partner notification. No log sources or event IDs apply — this is an information validation step, not a system-side forensic action.

Step 2: Do not redistribute — avoid forwarding this claim internally or to leadership as confirmed regulatory news; it has not been verified and may be fabricated

NIST Phase: Containment

Reference: NIST 800-61r3 §3.3 — Containment Strategy: limiting the spread and impact of an incident, applied here to the organizational spread of a potentially fabricated regulatory claim that could trigger premature or misdirected compliance spend

Controls: NIST IR-4 (Incident Handling) — containment scope includes preventing internal escalation of unverified threat claims that could cause operational disruption, NIST IR-6 (Incident Reporting) — reporting obligations require verified incident criteria to be met before notifying leadership or external stakeholders, NIST IR-8 (Incident Response Plan) — the IR plan should define communication gating: who may communicate what, at what verification threshold, CIS 7.2 (Establish and Maintain a Remediation Process) — a risk-based remediation strategy requires verified threat inputs; acting on fabricated claims wastes remediation capacity and may trigger false compliance obligations

Compensating: Issue a brief internal hold notice via your existing communication channel (email or ticketing system) to anyone who received the claim: 'This directive has not been verified against cisa.gov/emergency-directives — hold distribution pending confirmation.' Draft a one-paragraph holding statement for leadership that acknowledges the claim is circulating, states verification is in progress, and commits to a follow-up within 24 hours. No tooling required — this is a communication control action executable by one analyst.

Evidence: Document the internal distribution map of the claim before issuing the hold: who sent it, who received it, what channels were used (email thread headers, Slack channel names, ticket IDs). If the claim arrived via email, preserve the full email header (From, Reply-To, Received chain) to assess whether the source is a known threat intelligence vendor, an impersonator, or an unknown external party. This distribution record is relevant if the fabricated directive is later determined to be a deliberate influence operation targeting your sector.

Step 3: Assess AI surveillance exposure on its merits — inventory AI-driven video analytics systems in your physical security stack regardless of this claim, as adversarial ML risk to computer vision is a documented threat class

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: building and maintaining the defensive posture and asset visibility required to respond effectively to incidents, including identifying systems exposed to documented threat classes before incidents occur

Controls: NIST RA-3 (Risk Assessment) — requires assessing risk to organizational operations from threats relevant to the system, including adversarial ML evasion against computer vision, which is documented in MITRE ATLAS, NIST CM-8 (System Component Inventory) — requires maintaining an inventory of system components, including AI-driven video analytics platforms (e.g., Avigilon, Milestone XProtect, Genetec Security Center, Verkada, custom OpenCV-based pipelines) with their model versions and inference endpoints, NIST SA-9 (External System Services) — many AI video analytics systems rely on cloud-hosted inference APIs (e.g., AWS Rekognition, Azure Video Analyzer, Google Video Intelligence); these external dependencies must be inventoried as part of exposure assessment, CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory) — AI video analytics nodes, NVRs with embedded ML inference, and edge cameras with onboard AI (e.g., AXIS ACAP apps) must appear in the asset inventory with model firmware versions, CIS 2.1 (Establish and Maintain a Software Inventory) — AI model files (.onnx, .pb, .pt, TensorFlow SavedModel), inference runtime versions (TensorRT, OpenVINO, ONNX Runtime), and analytics software versions must be inventoried to assess patch surface

Compensating: Run a network scan using nmap against your physical security VLAN to identify camera endpoints and NVR/VMS hosts: 'nmap -sV -p 80,443,554,8080,8443 '. Cross-reference discovered hosts against your CMDB or a manually maintained spreadsheet. For each AI-enabled system, document: vendor, model, firmware version, inference type (on-device vs. cloud API), and network exposure (internet-facing vs. isolated). A 2-person team can complete this inventory for a mid-size deployment in one business day using nmap and a spreadsheet.

Evidence: Before beginning the inventory, capture the current network baseline for the physical security VLAN: export DHCP lease tables from your network infrastructure to identify all active endpoints, pull ARP tables from the access switch serving camera VLANs ('show arp' on Cisco IOS or 'ip neigh show' on Linux-based switches), and document any existing firewall or ACL rules governing camera-to-internet or camera-to-cloud traffic. This baseline establishes the

pre-assessment state and will be used to validate inventory completeness.

Step 4: Reference MITRE ATLAS — review adversarial ML evasion techniques relevant to image classification and object detection systems using the ATLAS framework at atlas.mitre.org

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: developing threat models and detection hypotheses based on documented adversary techniques to improve pre-incident defensive posture

Controls: NIST RA-3 (Risk Assessment) — adversarial ML evasion against computer vision is a documented threat class in MITRE ATLAS; risk assessment must account for techniques such as AML.T0015 (Evade ML Model) and AML.T0043 (Craft Adversarial Data) applied to video object detection, NIST SI-4 (System Monitoring) — monitoring requirements for AI video analytics should be informed by ATLAS technique patterns: specifically, detecting anomalous model confidence score distributions or unexpected classification label outputs that may indicate evasion attempts, NIST SA-11 (Developer Testing and Evaluation) — for internally deployed or customized computer vision models, adversarial robustness testing (e.g., using Foolbox, ART — Adversarial Robustness Toolbox) should be part of the model evaluation process, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — vulnerability management for AI systems must include tracking adversarial ML technique disclosures in MITRE ATLAS as a parallel to CVE tracking for traditional software, CIS 4.6 (Securely Manage Enterprise Assets and Software) — AI model files and inference runtime configurations should be version-controlled and integrity-verified to detect unauthorized model substitution, a precursor to adversarial manipulation

Compensating: Download the MITRE ATLAS Navigator (available as a standalone web app at https://mitre-atlas.github.io/atlas-navigator/) and filter techniques by tactic 'ML Attack Staging' and 'ML Model Access' to identify evasion paths relevant to your specific camera AI stack. For each AI video analytics vendor in your inventory, search ATLAS case studies for documented real-world incidents involving similar computer vision systems. Document applicable techniques (e.g., AML.T0015, AML.T0018 — Backdoor ML Model) as hunting hypotheses. No SIEM required — this is a manual threat modeling exercise.

Evidence: Before conducting the ATLAS review, collect baseline operational telemetry from your AI video analytics platform to establish normal model behavior: export model inference logs showing object classification labels, confidence scores, and detection frequencies for a representative 7-day period (log location varies by vendor — e.g., Milestone XProtect stores analytics events in its SQL event database; Genetec logs to its reporting module; custom pipelines may log to syslog or application logs). This baseline is required to detect future anomalies indicative of adversarial evasion (e.g., sudden drop in person-detection confidence scores, unusual label distribution shifts).

Step 5: Monitor CISA channels — subscribe to CISA's official directive and alert feeds to receive confirmed regulatory actions at the source, reducing exposure to fabricated or misattributed claims

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: updating processes and controls based on lessons learned to prevent recurrence, specifically hardening the organization's threat intelligence intake process to filter fabricated regulatory claims at the source

Controls: NIST SI-5 (Security Alerts, Advisories, and Directives) — requires receiving security alerts and directives from defined authoritative external organizations on an ongoing basis; CISA's RSS feeds (https://www.cisa.gov/cybersecurity-advisories/all.xml) and GovDelivery subscription service are the authoritative channels for Emergency Directives, NIST IR-8 (Incident Response Plan) — the IR plan should be updated to include a threat intelligence vetting procedure that routes unverified regulatory claims through a verification gate before internal distribution, informed by this incident, NIST IR-5 (Incident Monitoring) — tracking and documenting this fabricated directive claim as an incident record creates an organizational precedent for handling similar disinformation events in the future, CIS 8.2 (Collect Audit Logs) — audit logging for threat intelligence intake (who received the claim, what channel, what actions were taken) supports post-incident analysis and demonstrates due diligence if questioned by auditors about compliance with fabricated directives

Compensating: Subscribe to CISA's free GovDelivery email alert service for Emergency Directives and Known Exploited Vulnerabilities at https://public.govdelivery.com/accounts/USDHSCISA/subscriber/new — select 'Emergency Directives' and 'Cybersecurity Advisories' topics. Add CISA's RSS feed (https://www.cisa.gov/cybersecurity-advisories/all.xml) to a free RSS reader (e.g., Feedly free tier) configured to alert

on keyword 'Emergency Directive'. Document both subscriptions in your threat intelligence source registry. Total setup time: under 30 minutes for a 2-person team, no cost.

Evidence: As a lessons-learned artifact from this event, document the full lifecycle of this fabricated directive claim in your incident tracking system: origin source, initial receipt timestamp, internal distribution scope, verification steps taken, time-to-verification, and final disposition. This record supports NIST IR-5 (Incident Monitoring) documentation requirements and provides a reference case for training staff to recognize fabricated regulatory claims. Preserve the original claim artifact (screenshot, email, or URL) alongside this record for future reference if similar disinformation campaigns targeting AI surveillance compliance emerge.

Detection Guidance

No detection guidance can be issued for 'pixel blindness' as a named technique, it has no verified definition, published indicators, or documented attack chain. Issuing detection rules against an unverified technique risks wasting analyst resources and introducing false signal.

For organizations that independently choose to assess adversarial ML risk to their AI surveillance systems, the following areas are relevant based on documented adversarial ML research:

- Review AI video analytics system logs for anomalous confidence score distributions, consistently low detection confidence in specific camera zones may indicate model evasion rather than camera failure

- Audit physical security camera coverage for dead zones not captured by AI analytics processing, particularly in high-value access areas

- Test AI detection models against known adversarial input types (IR saturation, pattern-based clothing interference) if your vendor supports adversarial robustness testing

- Cross-reference AI-flagged and AI-missed events against raw footage to identify systematic detection gaps

These steps are grounded in documented adversarial ML research and represent sound security hygiene for AI-reliant physical security programs, independent of any claimed directive.

Framework Mappings

NIST-800-53R5

- **SI-4** — System Monitoring

CIS-V8

- **8.2** — Collect Audit Logs

NIST-CSF-2

- **DE.CM-01** — Networks and network services are monitored

Sources

Source	URL	Tier
How AI-powered video surveillance is changing security in 2026	https://sirixmonitoring.com/blog/ai-powered-video-surveillance-for-...	T3
AI Surveillance System Trends You Can't Ignore in 2025 - VORTEX	https://www.vortexcloud.com/resource/blog/surveillance-system	T3
AI-Powered Video Surveillance Enables Smarter Security Monitoring	https://www.sdmmag.com/articles/105258-ai-powered-video-surveillanc...	T3
The Complete Guide to AI Video Surveillance: Transforming Security ...	https://volt.ai/the-complete-guide-to-ai-video-surveillance-transfo...	T3
Preventing Slip and Fall Fraud with AI-Driven Security Camera ...	https://pavion.com/resource/preventing-slip-and-fall-fraud-with-ai-...	T3

Generated 2026-04-04 13:38 UTC by TJS Security Command Center