

Rethinking the Adoption of AI Surveillance Infrastructure:

An Information Systems Governance Failure

By: Gregory Nuter

Introduction

Automated License Plate Reader (ALPR) systems have been deployed across thousands of U.S. jurisdictions, with vendors such as Flock Safety now operating approximately 80,000 networked cameras that continuously capture images of vehicles and associated location metadata. Some of the basic data stored with these vehicle captures includes an AI-derived license plate number, time of capture, location, and vehicle characteristics (e.g., make, model, color, unique dents or scratches, etc.). These cameras' indiscriminate collection of vehicle information create highly sensitive records of individuals' daily lives. With a large network, law enforcement can gain a longitudinal view of an individual's life – from home and work locations to religious, political, and social activities. These surveillance systems, however, are frequently adopted without robust information systems governance, which includes comprehensive policies on data retention, access control, and auditing. Documented cases of abuse, misconfiguration, and security vulnerabilities have exposed a systemic problem with their adoption. Common incidents include law enforcement personnel misusing ALPR data to stalk former partners leaving poor reasons for searches such as 'investigation' (Kilmer, 2026). Other times, internet-exposed cameras were found indexed on platforms such as Shodan for anyone to view. This ease of exposure combined with large queryable datasets presents serious concerns for the privacy and security of the public. Several cities and counties have subsequently terminated or renegotiated ALPR contracts after uncovering security, auditing, or policy violations, revealing governance failures with tangible legal, operational, and reputational consequences. These outcomes are not isolated incidents.

In a time when our lives continually grow more digital, whether by choice or out of necessity, it is more important now than ever to advocate for data governance controls. Without clear policies on data retention, access control, auditing, and inter-agency data sharing, organizations utilizing these ALPR systems create avoidable security, legal, and ethical liabilities. Everyday citizens are the largest group who face daily risks when conducting their lives in cities and towns where these systems are

configured. For example, law enforcement officers are given full access to the ALPR databases in order to conduct investigations. This comes with the ability to search for specific license plates or types of vehicles. They are also given the ability to add individuals to a 'hotlist' to be alerted when the individual's vehicle gets detected by an ALPR camera. These actions are permitted without probable cause, a search warrant, or a need to attach a case number. Additionally, default configurations for access logs prove to be inefficient to determine intent or validity of searches. Through nationwide datasets, many officers default to using generic terms like 'inv' or 'investigation' to warrant their use of the database. This accumulates in a massive network of cameras tracking individuals' locations with metadata accessible under limited scrutiny. Various abuse cases arise from the nature of the data, its storage, and query-ability. One common abuse case is law enforcement officers utilizing the queryable database to track the movements of former relationship partners while going largely undetected. Another threat which arises from these systems is the misconfiguration of these ALPR cameras. Due to their connectivity to a network, misconfiguration has made some Flock Safety cameras accessible via websites like Shodan. Alongside misconfiguration concerns, documented, reviewed, and published security vulnerabilities draw concern around national security and whether adversarial nations and cyber-terrorist groups could utilize these cameras to spy on government officials, military personnel, military bases.

This paper will analyze these risks and failures within the broader domain of information systems governance, specifically addressing how inadequate frameworks in AI surveillance deployments lead to specific preventable risks such as privacy breaches and data misuse. This analysis will provide municipal CIOs, IT Directors, and city technology decision-makers with a pragmatic roadmap for aligning ALPR adoption with established standards. While there is a large number of ALPR vendors with vast product lines and features, this paper will focus primarily on Flock Safety, their ALPR hardware, and software services. The implications of this paper are not limited to Flock Safety and their

ALPR systems since governance failures and risk is a broader issue within the field of crime-fighting technology. Flock Safety has proven to be a useful example for discussion given their massive expansion. Importantly, the intent of this paper is neither to criticize the use of Automated License Plate Reader technology for criminal investigations nor the Flock Safety company.

Related Works

The existing literature on AI surveillance systems, particularly Automated License Plate Readers (ALPRs), is limited due in part to the recency of the technology and associated issues. However, earlier waves of networked crime-control technologies, such as wiretapping, predictive policing, and hotspot algorithms, have exposed a recurring pattern: technical capabilities expand faster than legal and governance safeguards. Historical analyses of U.S. electronic surveillance law show how statutory frameworks evolved in a piecemeal manner (e.g., wiretap prohibitions like Title III and FISA) repeatedly struggling to keep pace with ever-changing data interception methods and leaving gaps around stored communications and online monitoring (Young, 2011). Young also emphasizes that electronic surveillance routinely captures data beyond the scope of its initial investigative purposes and argues that core statutes like the Electronic Communications Privacy Act (ECPA) and the Foreign Intelligence Surveillance Act (FISA) remain insufficient and inconsistent, highlighting the lag of accountability structures. More recent work on predictive policing has documented how risk scores and hotspot algorithms were layered into existing law enforcement infrastructure with limited transparency and little oversight (Brayne, 2017). This historical narrative appears to repeat itself with modern policing technologies, particularly rapid ALPR adoption and deployments.

Literature on AI surveillance systems, particularly ALPR systems, remains sparse due to the newness of these deployments. However, existing research highlights critical governance and data management gaps that expose communities to privacy invasions, security breaches, and misuse. AI governance research broadly agrees that frameworks like COBIT 2019 and ISO 27001:2022 are essential to

mitigating risks in these data-driven systems. A systematic review of AI governance maps responsibilities, lifecycle stages, and accountability mechanisms, advocating for organization-wide approaches to ethical and compliant AI deployments (Batool et al., 2025). Complementary work emphasizes core data governance strategies, including access controls, encryption, explainable AI, and GDPR/CCPA compliance processes for secure AI-driven operations (Onoja et al., 2021). Although there is broad agreement on the value of these standards, there appears to be debate whether or not retrofitting legacy or rapidly deployed systems, where the speed of building infrastructure over enabling structured oversight, will sufficiently mitigate risk.

Security research on Flock Safety's hardware ecosystem demonstrates the consequences of weak governance practices at a service and hardware provider level. An exhaustive independent audit discovered 51 flaws across various Flock Safety systems, including gunshot detectors, ALPRs, and compute boxes. These vulnerabilities included disabled secure boot, unencrypted flash storage, and unlocked boot-loaders enabling root access (Gaines, 2025). Policymaker attention has also amplified these concerns. A federal letter to the Federal Trade Commission criticizes optional Multi-Factor Authentication (MFA), SMS-based logins, and password-sharing practices as factors that increase the risk of unauthorized access to ALPR data (Wyden & Krishnamoorthi, 2025). While vendors like Flock Safety assert that patches and configuration changes have been deployed, the lack of independent and/or field-wide validation raises questions about persisting risk in production environments.

Beyond technical vulnerabilities, critics and advocates have documented how ALPRs generate longitudinal "life maps" of vehicle movements that can be repurposed for surveillance of individuals rather than investigations of specific crimes. Analysis of Freedom Of Information Act released access-logs shows high volumes of queries justified with vague reasonings such as "investigation" often without associated case numbers or warrants, making meaningful auditing and accountability extremely difficult. This misuse risk aligns with prior reports of officers abusing ALPR systems to monitor former

partners, illustrating how weak governance directly translates into concrete harms. Policy work commonly frames these issues within a persistent crime-privacy tension, specifically statutes like California’s SB-34 and Virginia’s §2.2-5517 require ALPR policies and audits, yet researchers argue that gaps in enforcement and implementation undermine their effectiveness (Fash, 2019). Consequently, formal requirements for documentation, auditing, and retention limits often coexist with practices that still enable opaque, large-scale tracking of residents.

Empirical evaluations of ALPR effectiveness complicate the governance debate. Snow and Charpentier, Flock Safety employees, performed a study which reported an association between ALPR adoption and improved clearance metrics for certain criminal offenses; however, an independent case study also suggests that only a fractional number of license plate “hits” generate investigative leads (Hofer, 2022). This contesting evidence means municipalities could be accepting substantial surveillance infrastructure, and its associated risks, without a clearly established public-safety benefit.

Taken together, the literature reveals a consistent pattern: ALPR systems are being deployed without municipalities establishing governance mechanisms to secure sensitive data and ensure accountability. This gap between deployment scale and governance maturity motivates the present analysis.

Data Collection

Overview and Disclaimer

This analysis utilizes access logs procured through Freedom of Information Act requests to law enforcement agencies. It is neither meant to be treated as a representation of all agencies using Flock Safety, like-systems, nor a representation of to-date practices of these agencies. The analysis includes organizational and network access logs from 17 different organizations utilizing Flock Safety’s ALPR systems. This data was collected from these organizations via state-level Freedom of Information Act requests from 2024 to 2026, with majority coming from sources located within California, Washington,

and Illinois (see Appendix A). Majority sources of these audit logs comes directly from FOIA requests on MuckRock, a U.S.-based non-profit to help people access government information. Due to varying legislation, the released and withheld information can vary significantly. For example, some access logs list license plates and full names of queried suspects, while others exclude such data. Additionally, due to some prospective datasets not having enough un-redacted data or non-null reason columns (>50% empty, simple majority), they have been withheld from analysis. The determination to exclude datasets using a simple majority for redacted and null percentages of total was determined by calculating and plotting each datasets percentage composition of each respectively and plotting the distribution (See Appendix B).

Dataset

The dataset comprises approximately 20.1 million unique rows of access logs spanning 184 files. The data is stored in comma-separated values (CSV) format, recording officer name (full name, abbreviated, unique ID), queries against the ALPR network, and additional metadata such as time of query, search window, reason description, case number, filters, and organization (see Table 1). Each file was ingested with Polars, column names were standardized to lowercase and stripped of whitespace, and the tables were concatenated into a single dataframe capturing all recorded searches across agencies in a lossless join. To prepare the data for governance-focused analysis, the reason field was normalized by trimming extraneous characters and converting all entries to lowercase.

Column Name	Meaning	Example
user_id	Searching officer’s name, initials, or user id	“John Smith”, “J. Smith”, “J.S.”, “21452141”
Reason	Reason for querying system	“Hit and run”
Org name	Name of Officer’s organization	“Wheaton Police Department”
case_number	Case number associated with search	“WHP252164221”
Search time	Time querying was executed at	9/11/2024, 09:00:00 AM
Filters	Additional layer of filtering	“Washington”

Time frame	Window of time to search within	9/2/2024, 09:00:00 AM UTC to 9/24/2024, 12:00:40 PM UTC
Total networks searched	Number of networks searched in system	1,242
Total devices searched	Number of devices searched in system	4,215
Text prompt	Descriptive prompting for query	“Gray truck with mismatched wheels”
Moderation	Possible form of automatic moderation on queries.	Allow, warn, block

Table 1: Dataset Column Definitions

The goal of this exploratory analysis was to identify trends or potential areas for improvement by agencies. An area of focus for the analysis was to categorize the reasons officers stated they were using the system. These categories were researcher-defined and are not a legal determination. Through flat-joining regular expression patterns by group, logs were categorized into the following groups:

Category Name	Definition	Example
Legitimate	Reason clearly states the crime or purpose of query. Alleged-crime is clearly defined; legitimate categorization does not consider whether a case number is attached to the search	“hit and run”
Ambiguous	Reason is provided but is a non-obvious acronym or is not specific enough to an alleged crime. Often samples are categorized here because they are too broad and could be interpreted multiple ways.	“investigation”, “felony”
Concerning	Any logs which may raise concerns of misuse without further context. Within additional context, search could be explained away as legitimate however the current state could be cause for concern of misuse. Concern is further escalated when case number is not attached to the search.	“protest”, “pattern of life”, “threat”
Test	Any logs referencing ‘test’.	“focus test”, “test”
Numeric Code	Numeric value that cannot confidently be assigned as a case number. Criteria for this category is any logs which don’t meet ‘case number’ categories regular expression pattern	“99”
Case Number	Case number was filled into ‘reason’ column. Regex matches via {agency-prefix}{year}{case-number-suffix}. *Future analysis with additional agency data requires refining of	“WHP252164221”

	regular expression	
Null	Logs where ‘reason’ is null / empty string	
Redacted	Reason was redacted by agency before releasing logs	REDACTED
Uncategorized	Logs which don’t fit to current filtering.	“dhetf”, “file 1”

Table 2: Category Definitions

Filter lists were generated through trial and error, taking samples of the dataset and putting reason codes into a respective filter category to minimize the “uncategorized” entries. A second layer of filtering quantified how many of the ambiguous, concerning, and uncategorized logs were missing a case number. This allowed for a granular inspection of standard practice within these agencies and what gaps in current governance frameworks may exist.

Findings

Exploratory Analysis of FOIA Audit Logs

Table 3 shows that most ALPR searches are justified with vague justifications, missing case numbers, or both, which severely undermines audit-ability and governance. Across all categories, the majority of searches are lacking case numbers, ranging from approximately 78 percent of “legitimate” reasons to nearly all entries (94.56%) for entries with non-descriptive, non-case number numeric code.

Ambiguous, numeric-code, test, concerning, redacted, and uncategorized reasons together make up a large share of all queries (43.5%) while also having a lack of case numbers present ranging from approximately 88 to 97% coverage at an average of 93.92% lacking a case number regardless of categorization.

Reason Category	Count of searches	Missing Case Number	Percent of all searches (%)	Percent missing case # (%)
Legitimate	9,389,881	7,380,055	46.71	78.6
Ambiguous	6,610,373	5,983,339	32.88	90.51
Case Number	1,952,498	1,826,460	9.71	93.54

Uncategorized	1,128,491	1,000,777	5.61	88.68
Numeric code	906,374	857,088	4.51	94.56
Test	71,050	65,929	0.35	92.79
Concerning	33,185	33,185	0.17	100
Null	10,139	1,189	0.05	11.73
Redacted	33	32	< 0.1	96.97

Table 3: Prevalence of Vague Access Reasons




Table 4: Distribution of Searches by Reason Category

Additionally, total of 75.2 percent of searches are made without a case reference whether extracted from the dedicated ‘case_number’ column or through regular expression pattern matching from the reason column. Alongside this we observed that 0.01% or total of 1,975 undocumented searches occurred meaning there is no case #, reason, or text prompt. These searches require further exploration however as their minuscule presence possibly suggests unmarked redactions or the removal of sensitive data for purpose of FOIA release.

Privacy Concerns and Data Governance Failures

ALPR systems function as “dragnet” surveillance tools, meaning they capture data on all vehicles within their field of vision, regardless of whether the occupants are under suspicion of criminal activity (Brayne, 2017). This categorizes these systems as “bulk data collection tools”. This indiscriminate collection, combined with length of captured data and a large network of cameras, enables the creation of longitudinal “life maps” that reveal sensitive details of individuals’ daily activities, including home locations, workplaces, and social or religious affiliations bringing into question concerns over privacy (Brayne, 2017). This deeply personal tracking, often referred to as a 'pattern of life' investigation, will be further examined in the section analyzing search logs. While some states, such as California and Virginia, have implemented statutory requirements for ALPR policies, significant discrepancies persist between de jure regulations and de facto operational practices (Fash, 2019). For instance, evidence suggests that officers frequently access sensitive databases without probable cause, search warrant, or case number, reducing traceability. Although not illegal, depending on state law or agency policy, non-compliance risks arise as justification data remains optional. Furthermore, while ISO 27001 A.18.1.4 mandates data minimization and strict retention schedules, some jurisdictions have been exposed to maintaining data beyond policy-mandated retention limits.

Security Vulnerabilities and Technical Controls

Independent security research into the Flock Safety ecosystem has identified over 51 vulnerabilities that expose critical national security risks and reveal an absolute failure in secure development practices (Gaines, 2025). Although Flock Safety is used as an example given the expansive research by Jon Gaines of GainSec, it is notable to mention they are not the only company in the surveillance industry to deploy insecure technology into production environments (National Institute of Standards and Technology, n.d.). Also, Gaines’ research consists of vulnerabilities across Flock Safety’s hardware ecosystem, including devices like the Raven Gunshot Detection System, the Falcon, Sparrow, and Flex license plate readers, and Picard/Bravo compute boxes. These flaws included disabled secure boot,

unencrypted flash storage, and unlocked boot loaders that permit root access to the devices (Gaines, 2025). Such vulnerabilities enable the full compromise of the drive trust chain, potentially allowing adversarial actors to modify firmware or exfiltrate data. Furthermore, improper device configurations have resulted in internet-exposed cameras being indexed on public platforms like Shodan, rendering them accessible for unauthorized viewing (Gaines, 2026). Governance failures are also evident in user access management. The absence of mandatory multi-factor authentication (MFA) and the use of SMS-based logins increase the risk of unauthorized access by hackers or foreign intelligence services (Wyden & Krishnamoorthi, 2025). This lack of Multi-Factor Authentication (MFA) increases the risk of password sharing across agencies, a practice that has been documented in several cases, potentially compromising data security. This lax approach to authentication and security calls into question the adherence to governance by vendors as well as their customers employing their surveillance technology.

Misuse and Access Control Failures

Weak identity and access management governance has enabled systematic misuse of ALPR data for purposes unrelated to legitimate law enforcement activity. Multiple documented incidents expose personnel utilizing queryable databases to stalk former relationship partners while remaining undetected largely in part due to poor audit-ability (Kilmer, 2026). Access logs often reveal a high volume of queries justified by vague terminology such as “investigation”, which lacks the rigor required by COBIT 2019 DSS05.04 regarding the principle of least privilege. The results of these access logs are analyzed and explained further in a dedicated section later on. Freedom of Information Act request text messages also have revealed cases of password sharing between local police and DEA agents permitting unauthorized access and searching on the network (Henry, 2025).

Legal and Compliance Gaps

Despite legislative attempts to regulate ALPR use through statutes like California's SB-34, municipalities often lack the enforcement mechanisms required to ensure operational compliance (Fash, 2019). Additionally, these laws pose a baseline for agency policy, which is often broad and inadequate in establishing strong controls over data, often requiring law enforcement to establish an agency policy and make it publicly accessible. Many agencies rely on self-certification and vendor assurances rather than independent, third-party audits, which ISO 27001 A.18.2 identifies as essential for verifying compliance with security standards. The resulting governance gap is frequently discovered after the fact when security breaches or policy violations have occurred. Furthermore, vendor contracts often lack enforceable provisions regarding data privacy and security, leaving municipalities with significant legal and reputational exposure.

Interagency Data Sharing and Function Creep

The integration of ALPR data with regional fusion centers and federal agencies has created a "surveillant assemblage" where data collected for one purpose is repurposed for unintended activities, a phenomenon known as function creep (Brayne, 2017). Platforms like Palantir operate by integrating disparate data sources from the Department of Homeland Security (DHS), the Federal Bureau of Investigation (FBI), and Immigration and Customs Enforcement (ICE), facilitating the spread of surveillance into non-criminal justice institutions (Brayne, 2017). This loss of local oversight means that once municipal data enters a federal system, it may be used for previously non-consented or policy-violating purposes. Current government frameworks lack standardized interagency sharing agreements that address security, privacy, and accountability tracking across multiple jurisdictions.

AI-Specific Governance Challenges

ALPR systems, as AI-driven technologies, introduce unique challenges related to algorithmic transparency and bias mitigation that traditional IT governance frameworks are not equipped to address (Onoja et al., 2021). Most ALPR vendors do not disclose plate recognition accuracy rates or false

positive rates, creating a 'block box' that prevents independent validation of system reliability (Brayne, 2017). Additionally, there is a lack of documentation regarding training datasets used for these AI models, raising significant concerns about representativeness and potential for algorithmic bias. Although these concerns remain low for AI-based license plate detection, they are raised as function creep takes hold and aggregation services like Flock Nova or Palantir utilize these vast datasets to find patterns and relationships not previously discoverable using standard analysis methods. Systematic literature reviews show only a small fraction of studies address the “who”, “what”, “when”, and “how” of AI governance (Batool et al., 2025), leaving municipal decision-makers without a clear roadmap for ethical deployment. This results in systems being deployed without completing critical lifecycle stages, such as ethical reviews or performance monitoring (Batool et al., 2025).

Governance Gaps

The analysis of current Automated License Plate Reader (ALPR) deployments indicates that technology adoption consistently outpaces the implementation of governance frameworks, creating risk exposure. This risk exposure affects all parties involved, from the community to law enforcement agencies and nation (Brayne, 2017). This phenomenon follows a pattern observed in earlier crime-control technologies, where technical capabilities expand faster than legal and governance safeguards (Young, 2011). Adopting institutions seem to prioritize the perceived benefits of surveillance infrastructure without establishing and enforcing necessary governance mechanisms, such as those defined in ISO/IEC 27001 or COBIT 2019. These frameworks require risk assessments prior to system deployment in an attempt to minimize risks – which is crucial for highly sensitive systems such as surveillance networks. Consequently, many ALPR systems operate within a "lag" period, characterized by a lack of transparency and oversight (Brayne, 2017). Notably, after adoption, agencies do not enforce strict policies on data and access controls due to loosely defined policy and a lack of meaningful legislation at the state level.

ALPR database search practices in the sampled agencies are best understood as a misalignment with governance controls. Across millions of queries, officers frequently rely on vague justifications or ambiguous labels, with a substantial majority lacking a case number. This pattern is misaligned with COBIT 2019 DSS05.04 (“Manage user identity and logical access”), which expects least-privilege access based on clearly defined business needs and sufficient logging and traceability so that each query of the system is justified and attributable (ISACA, 2022). Current analysis suggests traceability is nearly useless. Similarly, ISO/IEC 27001:2022 Annex A.9 (“Access control”) expects access rules grounded in documented authorization, while Annex A.12.4 (“Logging and monitoring”) and A.18.1 (“Compliance with legal and contractual requirements”) require event records that are rich enough to reconstruct security-relevant actions to demonstrate conformity to law and policy (International Organization for Standardization, 2022). In practice, the absent case identifiers produce formally complete but substantively weak audit trails that do not meet underlying objectives of the aforementioned standards.

This misalignment extends beyond application-level logging and into broader questions of lifecycle governance. AI governance literature emphasizes that high-risk AI systems handling sensitive personal or location data should be governed across all stages (pre-deployment risk assessment, design and build, and post-deployment monitoring) with clearly assigned accountability for both data and systems. ISO/IEC 27001 A.18.1.4 (“Privacy and protection of PII”) and A.18.2 (“Information security reviews”), reinforce this expectation by calling for data-minimizing processing and privacy-preserving controls. This also includes independent evaluations of control effectiveness, while COBIT 2019’s MEA01 (“Monitor, evaluate, and assess performance and conformance”) and MEA03 (“Monitor, evaluate, and assess compliance with external requirements”) require systematic assurance that security practices match organizational risk appetite and regulatory obligations (ISACA, 2022; International Organization for Standardization, 2022). Yet the rapid deployment of networked ALPRs, documented device-level

security flaws, reliance on optional or weak authentication, and the absence of independent technical audits indicate that municipalities and vendors are only partially realizing the intent of frameworks in day-to-day operations.

Discussion

The findings suggest that powerful surveillance systems are being deployed without proper policies for data governance and access. Paired with limited oversight and audit-ability, these systems remain widely accessible with minimal preventative measures for abuse. Evidence indicates that governance maturity is lagging behind deployment, leaving adopting municipalities exposed to liability without proper controls. This issue is significant because analysis suggests this pattern is repeated across departments nation-wide.

One key implication of the findings is that weak documentation practices are substantially undermining accountability. The audit log analysis shows that a large share of ALPR system queries were associated with vague reasons, missing case numbers, or both, and that 75.2 percent of searches lacked a case reference when combining the dedicated case number field with case-number pattern matching. In governance terms, this suggests that many agencies possess a perceived level of traceability through logging while failing to create conditions for meaningful audit-ability. Rather than functioning as a strong accountability mechanism, existing logs are substantively thin in content and context rather than rebuilding intent.

The findings also reshape how the problem of ALPR deployments should be understood. The issue not only lies in surveillance expansion, but expansion under administratively weak conditions. Findings repeatedly show the governance gap visible across multiple layers of the system, including user behavior, agency policy, vendor security practices, and interagency access arrangements. For that reason, evidence points toward a broader information systems governance failure rather than a narrowly technical security problem. Though these ALPR systems can serve legitimate public-safety

reasons for the adopting agencies, the utility alone does not offset the need for controls capable of constraining misuse and limiting the unnecessary accumulation of sensitive data.

The security findings further reinforce this sentiment. Governance failures not only arise from officer misuse, but also weaknesses in software and hardware. Vulnerabilities identified in the broader Flock Safety ecosystem suggest agencies may inherit risk from vendor ecosystems while also amplifying risk through local misconfiguration. In that sense, deployment without established and strong governance frameworks appears to shift municipalities into a reactive posture, exposing themselves greatly to risk rather than bolstering a system build up in a preventative manner.

Additionally, these findings pose concerns around organizational decision-making. ALPR governance cannot be treated as a narrow law-enforcement operations issue, because the risks cut across legal compliance, cybersecurity, records management, procurement, and public trust. This suggests that municipal CIOs, IT directors, and other decision-makers should view ALPR adoption as an enterprise governance question rather than as a stand-alone policing technology purchase. Framed this way, the problem becomes less about whether agencies can deploy the technology and more about whether they have the institutional capacity to govern it before it is integrated into routine practice.

Function creep introduces ever-increasing apprehensions regarding surveillance and the normalization of broad data collection and access. Because ALPR systems collect location-linked vehicle data at scale, weak controls over what data is collected can gradually expand use beyond its initial justification. The discussion of “pattern of life” searches suggest failures do not merely create accidental risk but also create conditions where expanded surveillance becomes operationally easy and administratively difficult to challenge. In these situations, the absence of proper regulatory governance before deployment could allow for the misuse of the technology for broader intelligence gathering or continuous surveillance.

Taken together, these findings suggest that municipal decision-makers may benefit from a governance-first roadmap for ALPR adoption. Rather than treating deployment as the starting point, organizations should require a documented use case, defined access and retention policies, independent security validation, and ongoing audit oversight before and after implementation.

Roadmap for Decision-Makers

For municipal CIOs, IT directors, and others responsible for approving and implementing these systems, a plan would start with defining what constitutes a proper use case, documenting the grounds upon which access was granted legally, and determining whether an agency has the ability to enforce the ability to retain data, to have data audited, and role-based access before going into operational mode. Agencies should then demand baseline security and responsibility protocols from vendors and internal IT teams. This involves the independent validation of device security, the enforcement of multi-factor authentication, query logging tied to case numbers where needed, access log analysis, and ensuring that any contract language allows for enforcement of security and privacy provisions.

The following might serve as the outline of such a roadmap:

Stage	Stage Name	Description
0	Prior to deployment	Establish the public safety purpose of the proposed implementation, assess privacy and security risks, and ensure that this is proportional to the problem at hand
1	Control development	Set guidelines for retention of data, access controls, auditing, and inter-agency sharing before the system is put into operation.
2	Technical requirements	Validate authentication, device security, patching practices, and logging independently.
3	Governance structure	Assign responsibility for ongoing oversight, audit, and compliance checking to maintain oversight.
4	Reevaluation	Periodically reevaluate the necessity, utility, and risks associated with the continued use of the system especially if the scope grows.

Table 5: Pragmatic Roadmap For Decision-Makers

Conclusion

Evidence suggests that ALPR adoption becomes most problematic when governance is treated as an afterthought rather than a prerequisite. This hypothesis however could be expanded to the greater field of surveillance infrastructure. In reviewing the evidence there appears to be a deeper organizational issue where sensitive surveillance systems are being adopted and deployed before the controls needed to manage them responsibly are in place. This shifts the main question for decision makers from whether these systems assist investigations to whether these institutions have the capacity to govern them in a manner which protects the public and preserves accountability. When speed of deployment is prioritized, risk is dramatically increased for privacy, trust, and legal compliance.

A better approach is to treat governance as a necessary part of deployment itself. Frameworks such as COBIT and ISO 27001 shift the attention toward access control, documented purpose, review-ability, retention limits, and security assurance before surveillance infrastructure become thoroughly ingrained in daily practice. Responsible adoption of ALPR systems depends less on the promise of the technology than on whether institutions are willing to establish meaningful oversight before the system goes live.

Works Cited

- Batool, A., Zowghi, D., & Bano, M. (2025). AI Governance: A systematic literature review. *AI and Ethics*, 5(3), 3265–3279. <https://doi.org/10.1007/s43681-024-00653-w>
- Brayne, S. “Big Data Surveillance: The Case of Policing.” *American Sociological Review*, U.S. National Library of Medicine, Oct. 2017, [pmc.ncbi.nlm.nih.gov/articles/PMC10846878/](https://pubmed.ncbi.nlm.nih.gov/articles/PMC10846878/).
- Business & Human Rights Resource Centre. (2025, August 26). United States: City of Evanston ends Flock Safety contract after illegal data sharing. <https://www.business-humanrights.org/en/latest-news/united-states-city-of-evanston-ends-flock-safety-contract-after-illegal-data-sharing-with-federal-agencies/>
- California State Legislature. (2023). SB 34: Automated license plate recognition systems: use of data. 2023–2024 Regular Session. <https://legiscan.com/CA/text/SB34/2023>
- City of Mountain View. (2026, February 25). City Council terminates ALPR contract with Flock Safety [Press release]. www.mountainview.gov/Home/Components/News/News/1211/284
- Fash, L. (2018). Automated License Plate Readers: The Difficult Balance of Solving Crime and Protecting Individual Privacy. *Md. L. Rev. Endnotes*
- Gaines, J (2025). GainSec/anti-crime-ecosystem-research: v1.2 Public Release – Examining the Security Posture of an Anti-Crime Ecosystem (v1.2-PR). Zenodo. <https://doi.org/10.5281/zenodo.17584876>
- Gaines, J. (2026). *Finding 67 flock safety live ptz camera/LPR feeds and debug web interfaces accidentally exposed without authentication to the internet*. GainSec. <https://gainsec.com/2026/01/09/bird-hunting-season-finding-67-live-camera-feeds-and-debug-web-interfaces-accidentally-exposed-by-flock-safety/>

Henry, C. (2025, August 28). After Oak Park cut ties, state says Flock Safety broke the law. Oak Parker. <https://www.oakpark.com/2025/08/28/state-says-flock-safety-broke-the-law/>

International Organization for Standardization. (2022). ISO/IEC 27001:2022 Information security, cybersecurity and privacy protection—Information security management systems—Requirements (3rd ed.). <https://www.iso.org/standard/27001>

ISACA. (2022). COBIT 2019 framework: Governance and management objectives (2nd ed.). ISACA. <https://www.isaca.org/resources/cobit>

Kilmer, G. (2026). Milwaukee cop used flock cameras to track romantic partner. Urban Milwaukee. <https://urbanmilwaukee.com/2026/02/25/milwaukee-cop-used-flock-cameras-to-track-romantic-partner/>

National Institute of Standards and Technology. (n.d.). NVD - CVE-2017-7921. National Vulnerability Database. <https://nvd.nist.gov/vuln/detail/CVE-2017-7921>

Onoja, J. P., Hamza, O., Collins, A., Chibunna, U. B., Eweja, A., & Daraojimba, A. I. (2021). Digital transformation and data governance: Strategies for regulatory compliance and secure AI-driven business operations. *J. Front. Multidiscip. Res*, 2(1), 43-55.

Snow, Adam, et al. "Flock Safety Technologies in Law Enforcement: An Initial Evaluation of Effectiveness in Aiding Police in Real-World Crime Clearance." *ResearchGate*, Jan. 2024, www.researchgate.net/profile/Johnny-Nhan/publication/377845222_Flock_Safety_Technologies_in_Law_Enforcement_An_Initial_Evaluation_of_Effectiveness_in_Aiding_Police_in_Real-World_Crime_Clearance/links/65c1020d7900745497673897/Flock-Safety-Technologies-in-Law-Enforcement-An-Initial-Evaluation-of-Effectiveness-in-Aiding-Police-in-Real-World-Crime-Clearance.pdf.

State v. Ayala (2026). Criminal complaint (DA Case No. 2026ML004796). Milwaukee County District Attorney's Office. https://mkepdpio.org/wp-content/uploads/2026/02/Criminal-Complaint_3-Ayala-Josue.pdf.

Virginia General Assembly. (2025). Use of automatic license plate recognition systems, § 2.2-5517, Code of Virginia.

Wyden, R., & Krishnamoorthi, R. (2025, November 3). Letter to FTC Chair Andrew Ferguson requesting investigation of Flock Safety cybersecurity practices [Letter]. United States Senate & House of Representatives.
https://www.wyden.senate.gov/imo/media/doc/wyden_letter_to_ftc_on_flockpdf.pdf

Young, Mark D. "Electronic Surveillance in an Era of Modern Technology and Evolving Threats to National Security." *Docslib*, 6 Apr. 2011, docslib.org/doc/13388735/electronic-surveillance-in-an-era-of-modern-technology-and-evolving-threats-to-national-security-mark-d.

Appendix


Appendix A

Dataset Makeup by Year and Agency

Organization	2026	2025	2024
Aurora, IL	Partial	Partial	
Boulder County Sheriff Office (BCSO)	Partial		
Benton County, WA		Partial	
Boulder, CO	Partial		
California Highway Patrol		Complete	
El Cerrito, CA	Partial		
Ellensburg, WA		Partial	
Lincoln, WA	Partial	Partial	
Lodi, CA			Partial
Lynwood, WA		Partial	
Niles, IL		Partial	
Pasadena, CA		Partial	Complete
Roselle, IL		Partial	
Spokane, WA		Partial	
Tonawanda, NY		Partial	
Wheaton, IL		Partial	Partial
Yakima, WA		Partial	

Appendix B

Distribution of Null and Redacted Value as Percent of Respective Dataset



Redacted Percent By Dataset

