

Praeva Analytics ALPR Audit Records Civic Review Packet

Evidence, accountability, and a path to investigation and removal

Prepared for	Subject	Review date
Residents, city staff, and elected officials	Ventura County Sheriff automated license plate reader records	September 2026

Purpose

This packet is designed to support public action. It explains the records request, the agency response, the analytical work completed from the delivered audit exports, and the questions officials must answer before this camera system is allowed to expand, renew, or continue without stronger safeguards. It is written so that residents can understand the core findings while data and policy reviewers can examine the methods and evidence limits.

The central conclusion is straightforward: the delivered records show a large volume of recorded searches, long requested time windows, meaningful changes over time, and months with unusually high concentrations of extreme window lengths. The incomplete audit trail prevents the public from determining who authorized those searches, why they were conducted, what they returned, or whether they complied with policy. That combination demands an independent investigation, an immediate pause on expansion and renewal decisions, and a defined path to remove cameras or network access that cannot be shown to be necessary, effective, and accountable.

Meaningful oversight now depends on complete row-level audit information. Legally required redactions should be targeted to specific protected fields rather than used to strip the records of the structure needed for accountability. If the agency and vendor cannot produce auditable records, that failure should weigh against continued deployment.

Prepared from the supplied VCSO public-records response, policy and procurement materials, monthly Organization Audit and Network Audit exports, and the supplemental visual-analysis series described in this packet.

How to use this packet

Readers do not need a statistics background to use this packet. Start with the executive summary and the four plain-language rules below. The visual narrative then moves from broad network context to local agency comparisons, search-window scope, population scale, and the specific records needed for follow-up.

When the packet says	It means	It does not mean
Audit-log row	One delivered system record attributed to an organization label	One unique person, vehicle, investigation, or resident
Requested window	The span between the query start and end timestamps	How long data were retained or how many results were returned
Networks or devices searched	Scope values reported by the export	People searched, matches returned, or the physical area searched
Extreme row	A duration above a defined statistical threshold	An improper, illegal, or unjustified search

Evidence labels used throughout

- Directly observed: a value or pattern present in the delivered records.
- Review signal: a pattern strong enough to justify record-level or policy review.
- Alternative explanation: a plausible non-misconduct cause that should be tested, such as an export change, reporting change, duplicate delivery, operational shift, or data-entry problem.
- Not established: a conclusion the available records cannot support.

Contents

1. Executive summary and requested discussion actions
2. The public-records request and agency response
3. Data foundation, quality controls, and interpretation rules
4. Visual narrative and explanation of each analysis output
5. The case for action and alternative explanations
6. Questions officials must answer and recommended next records
7. Technical methods and output index

Executive summary

The delivered material is substantial enough to identify patterns that require action, but incomplete for determining the purpose, user, case context, result, or geographic footprint of individual searches. The absence of those fields is not a reason to accept the program on trust. It is a reason to investigate before public officials authorize continued operation, expansion, or renewal.

Five findings that should guide the discussion

1. The export is large and complex. The controlled analytical baseline contains 5,265,993 delivered observation rows, including 4,823,632 rows eligible for the primary comparison period and 27 Network Audit months from January 2024 through March 2026. March 2026 is partial.
2. The Ventura County Sheriff label is associated with 129,109 delivered rows across 27 active months. The exact-distinct count is 128,021, showing that duplicate sensitivity matters but does not explain away the overall volume.
3. Long requested search windows are common enough to require policy review. For the Sheriff label, the median window was seven days, the 90th percentile was 2,019 hours, the 99th percentile was 8,762 hours, and the maximum was 201,888 hours, or about 23 years.
4. The long-window tail was not evenly distributed over time. Using a robust threshold of about 906 hours, six months in 2024 were at least 1.25 times the full-period extreme rate. January and September 2024 had the highest exposure-adjusted ratios.
5. The records do not show who was searched, where a request was aimed, whether a search returned results, or whether a particular search had a legitimate purpose. Those questions require the omitted or unavailable audit fields and supporting governance records.

Recommended immediate actions

Praeva Analytics recommends that the governing body direct the following actions now:

- Pause new camera deployment, contract expansion, and renewal commitments until the investigation is complete and the public receives a written necessity and proportionality finding.
- Request a complete row-level audit export with the legally disclosable portions of user, purpose, reason, and case or report identifiers preserved.
- Ask VCSO and the vendor for the export data dictionary, schema version, field definitions, timezone rules, duplicate-generation behavior, and a description of known changes between monthly exports.
- Require a written reconciliation of the six elevated 2024 months and the longest-duration rows, including whether repeated boundary values reflect policy defaults, user choices, migration artifacts, or errors.
- Document current training, authorization, recurring audit, retention, sharing, and exception-review practices even where earlier records were not maintained in a reportable form.
- Commission an independent technical and policy review rather than relying solely on the operator or vendor to evaluate its own system.
- Return with a public staff report that identifies which cameras and network privileges are demonstrably necessary and recommends removal, disconnection, or non-renewal where necessity, effectiveness, and accountability cannot be shown.

The request and what the agency provided

The Ventura County Sheriff's Office stated that it received the request at the Thousand Oaks and Moorpark stations on August 21, 2026, and issued its response on September 14, 2026 after taking an extension. The request was organized into three groups: audit and access records, governance and administration records, and deployment and procurement records.

Request area	What was requested	Agency response described in the letter	Oversight implication
Audit trails and access	System-generated searches, logins, access activity, statutory fields, search controls, audits, and external-network activity	Organization Audit and Network Audit logs were produced. The agency said this was the only available responsive information for most audit items.	The logs support volume, timing, scope, and duration analysis, but not a complete reconstruction of individual search purpose or result.
Reason and case context	Purpose or reason and case or report number fields	The response said the reason field was omitted under an investigatory-records exemption.	Purpose-level testing is not possible from the delivered analytical fields. A field-specific redaction approach would preserve more audit utility.
Moorpark user mapping	Records mapping Moorpark Station users to accounts	The response said this specific information was not tracked and no responsive record existed.	The analysis cannot attribute rows to individual users or isolate Moorpark users without unsupported inference.
Policy and oversight	Policy versions, oversight personnel, training, sharing, configuration, retention, errors, misuse, and audit exceptions	Policy 429 and an inter-agency MOU were referenced. The response identified shared oversight roles, said formal ALPR training records were not maintained, and stated that some configuration and retention records were unavailable or not tracked.	Written policy can be compared with observable logs, but implementation cannot be fully verified without recurring audit and training documentation.
Procurement and deployment	Contract, amendments, camera inventory, expansion, and Moorpark-specific deployment	Contract and grant materials were produced. The response said a readily exportable inventory and Moorpark-specific deployment records were not maintained. It described 147 cameras as the Sheriff's total, including contract cities.	The supplied records do not support a precise Moorpark device inventory or a camera-to-search attribution model.

Related policy and procurement context

Policy 429 states that ALPR use is restricted to official and legitimate law-enforcement business, access is limited to designated staff, users must complete department-approved training, stored data must be protected through login controls that document access by user, date, and time, and system audits should be conducted regularly. The 2022 policy also states a five-year server retention setting, subject to case-related exceptions.

The contract materials describe Flock network and search capabilities, including local and broader network access, time-and-location searches, license-plate lookup, vehicle-attribute search, audit reporting, and other platform functions. A supplied grant budget includes a \$997,000 line for 100 automated license plate reader camera services within a broader organized-retail-theft program. These documents show capability and funding context; they do not prove how any specific search was conducted.

Sources: VCSO response letter dated September 14, 2026; Policy 429 dated December 23, 2022; supplied Flock contract materials; supplied BSCC Organized Retail Theft Prevention Grant budget workbook.

Data foundation and quality controls

The analysis uses provenance-preserving rows from monthly exports. It keeps the delivered record population visible while also calculating exact-distinct measures so that duplicate delivery can be measured rather than silently erased. Candidate geography and organization classifications are kept separate from human-approved classifications.

Measure	Controlled value	Why it matters
Delivered observation rows	5,265,993	Shows the full quantity inducted from the delivered exports.
Primary-period eligible rows	4,823,632	Supports comparable analysis within the controlled period.
Monthly source-coverage rows	32	Tracks which monthly sources were available and whether a month was complete.
Network Audit months	27	Covers January 2024 through March 2026; March 2026 is partial.
May 2024 delivered rows	645,233	A major reconciliation month with 202,872 eligible rows and 442,361 source-period mismatches.
May 2024 cross-export exact duplicates	434,599	Demonstrates why duplicate and source-period sensitivity are required.
May 2024 residual discrepancy rows	7,762	Shows that duplicates account for most, but not all, of the mismatch population.

Quality principles

- Preserve raw evidence and source hashes; perform analysis in derived outputs.
- Report delivered and exact-distinct values together when duplicates can change interpretation.
- Treat missing, unavailable, absent, and numeric zero as different states.
- Keep reporting month, claimed source period, and query timestamps separate.
- Do not convert candidate geography into a factual claim without human approval.
- Publish coverage percentages for device-based measures and describe incomplete coverage as a lower bound.
- Retain temporal exceptions in observed totals, report them separately, and test sensitivity to their exclusion.

What the data can and cannot answer

Supported by the delivered logs	Not supported without additional records
How many rows were delivered and attributed to a label	How many unique people or vehicles were searched
The reported query-window start and end timestamps	Why a specific query was run or whether it was justified
Reported network and device scope for covered rows	Which physical cameras were searched in a specific request
Temporal patterns, duplicate sensitivity, and statistical tails	Whether results were returned, viewed, downloaded, or used
Differences among organization labels and time periods	A finding of illegality, abuse, overreach, or misconduct

Part One From the national network to the local question


The visual sequence begins outside Ventura County because local search records sit within a wider networked system. The first two graphics show that non-California labels vary sharply in frequency, persistence, and requested scope, and that a relatively small set of labels accounts for a large share of requested device-hours. The point is not to rank agencies as good or bad. It is to demonstrate why simple row counts are not enough: frequency, duration, reach, persistence, and duplicate structure can tell different stories.

Non California organization profile

Figure 1 compares prominent non-California organization labels across several dimensions. Its left panel ranks frequency. Its upper-right view contrasts typical and longer search windows among frequency leaders. Its review table identifies combinations that deserve closer inspection. The visual is a screening tool: labels may represent agencies, centers, task forces, or export naming conventions, and classification remains subject to reconciliation.

The practical lesson is that high frequency and long windows are separate behaviors. A label can appear often but use comparatively shorter windows, or appear less often while contributing unusually large requested scope. A responsible review therefore avoids a single-score accusation and instead asks which combination of measures is driving a label's position.

FIGURE 1



Requested device hour concentration

Figure 2 asks a different question: which labels account for the greatest total requested device-hours? Device-hours combine a reported device count with a requested duration. This helps identify where the broadest requested scope accumulates, but it is not a count of returned records and not a measure of people affected.


The chart shows concentration among a small set of labels. That concentration is useful for prioritizing review because a modest number of records can carry a large share of total scope when windows are long or device counts are high. The adjacent coverage check is essential: device-hour totals only represent rows where a usable device value was present. Missing device data are treated as unavailable, not as zero.

FIGURE 2

Requested Device-Hour Scope Is Concentrated Among a Small Set of Non-California Labels


Device-hours = devices_searched × query-window hours, summed across valid primary-period Network Audit rows.

Top 20 labels by total requested device-hours



TOP-20 TOTAL | 2707.52B requested device-hours. 20 of 20 geography classifications are candidate-unreviewed.

Coverage check



QUALIFICATION | Device-hours describe reported requested scope. They do not establish returned records, unique devices or people, retention, search direction, sharing, purpose, legality, or misconduct.

Part Two Placing the Ventura County Sheriff label in context

The next four creations test whether the Ventura County Sheriff label looks unusual under several comparison frames. No single peer group is perfect. County sheriffs are functionally intuitive, a volume-matched activity band reduces scale mismatch, a population band adds jurisdiction size, and a nationwide sheriff set adds breadth. The analysis therefore reports multiple frames and a sensitivity view instead of selecting the one that produces the strongest result.

Peer percentile profile


Figure 3 compares the Sheriff label across frequency, exact-distinct frequency, requested device-hours, longer-window percentiles, persistence, and typical reported reach. The broad pattern is that the label ranks high on frequency and long-window measures in several peer definitions. Reach measures are more mixed, reminding readers that a label can be highly active without ranking at the top on every scope field.

Percentiles describe relative position within a selected comparison group. They are not probability statements, performance grades, or findings of wrongdoing. Small cohorts are especially sensitive to membership decisions, which is why the visual names its peer frames and the companion analyses test alternative weightings.

FIGURE 3

Where VCSO Falls Across Multiple Peer Spectrums

Orange dashed line = peer median (50th percentile)



VCSO ranks high in frequency, long search windows, requested device-hours, and persistence across most peer groups. Its typical numbers of networks and devices are much lower, especially compared with the nationwide sheriff group.

Candidate-qualified comparison; percentiles describe the export population and are not findings of misconduct.


Monthly percentile trends

Figure 4 moves from a full-period ranking to month-by-month position among nationwide sheriff labels. Frequency and requested device-hour percentiles remain high through much of the period. P90 window, networks, and devices show more movement, including a visible break around mid-2024 and additional shifts in early 2025.

This matters because an annual average can hide a temporary expansion, reporting transition, or operational regime change. The chart does not identify the cause of those changes. It identifies when to look. The correct follow-up is to compare source files, export structure, policy settings, staffing, and record-level distributions on both sides of each boundary.

FIGURE 4

VCISO's Monthly Position Within the Nationwide Sheriff Peer Universe



VCISO stays well above most nationwide sheriff labels in monthly frequency, requested device-hours, and long search windows. Its long-window percentile remains consistently high even when the other scope measures move sharply. Reported scope shifts from broader device reach in early 2024 to fewer typical devices later, while typical network reach briefly rebounds during 2025.


Dotted segments bridge months with unavailable values; they are not measured points. March 2026 is partial, and percentiles do not establish directionality or misconduct.

Regime comparison

Figure 5 groups the timeline into interpretable periods: a recovered early-2024 period, a mid-2024 through January 2025 baseline, a February-March 2025 expansion, a post-expansion period, and partial March 2026. The heat map shows that frequency-related measures remain elevated across regimes while median device and network measures shift more substantially.

The purpose of the regime view is methodological discipline. It prevents a short expansion from being treated as the permanent baseline and keeps the partial final month from being compared as if it were complete. The analysis should be updated if a better documented boundary becomes available.

FIGURE 5




Composite sensitivity

Figure 6 tests whether the comparative conclusion depends on how measures are weighted. Equal weighting, requested-scope emphasis, and frequency emphasis produce different composite values, but VCSO remains above the peer midpoint in each displayed scenario. The population-band comparison is especially high, while the activity-band frame is comparatively more restrained.

Composite scores are useful for sensitivity testing, not for making a final judgment. Weighting is a value choice. For that reason, the packet keeps the underlying components visible and does not reduce the investigation to a single league table.

FIGURE 6

VCISO's Composite Position Is Tested Across Three Weighting Scenarios



VCISO ranks above the peer median under every weighting method and comparison group shown.

The population-based comparison places VCISO especially high, while the activity-based comparison produces the lowest scores.

Its position rises most when frequency and persistence receive more weight, showing that the overall result depends partly on what the score emphasizes.

Composite scores summarize observed metrics; metric-by-metric results remain controlling. Candidate-qualified and not evidence of misconduct.

Part Three Separating Ventura-area labels and agency roles

Export labels that contain the word Ventura do not necessarily describe the same agency, function, or jurisdiction. The next visual keeps Sheriff, District Attorney, municipal police, dispatch, and other labels separate. This avoids combining rows merely because names share a county reference.

Ventura area organization profile


Figure 7 compares Ventura-area labels on frequency, scope, persistence, and change. The Sheriff label has the largest exact-distinct requested device-hour total among the displayed Ventura-area labels and ranks high on multiple measures. The District Attorney label is less frequent but later analysis shows a substantial long-window tail.

The strongest use of this graphic is organizational hygiene. It makes clear that a county label cannot be assumed to be a countywide action and that the same export may contain labels for different operational units. Each label requires its own reconciliation before conclusions are made about governance or responsibility.

FIGURE 7

Ventura County-Area Organization Labels: Frequency, Scope, Persistence, and Change

Tie-aware percentile ranks compare only these 7 matched labels; 50 marks the middle of this small local cohort.



Selection uses a reviewable place-name roster for Ventura County cities and communities. A name match does not by itself establish geography, agency identity, search direction, returned records, purpose, legality, or misconduct.

Part Four What the search-window profiles reveal

The three search-scope profiles use a common structure: monthly frequency and requested scope, the full distribution of window lengths, typical monthly network and device reach where reported, and a carefully qualified comparison. Reading the profiles together helps separate what is common across the export from what is specific to one label.

Ventura County District Attorney


Figure 8 shows that the District Attorney label combines modest frequency with long search windows. This is an important example of why a low row count should not be treated as low scope. A smaller number of requests can still accumulate substantial requested device-hours when time windows are long or reported reach is broad.

The figure is descriptive. It does not establish the subject of a search, the user, the investigation, or whether the request returned any records. It does identify a review priority: long-window rows should be reconciled to purpose, case context, policy, and any system default that could explain repeated durations.

FIGURE 8

Ventura County District Attorney: Long Search Windows Despite Modest Frequency

1,859 attributed rows across 24 active months • P90 window 2,210 hours • observed exact-distinct scope 8,396,477,528.21 device-hours (79.2% device-row coverage)



This label ranks low in local frequency but highest in P90 search-window length, so each request may cover a broader historical period than the count alone suggests. Use of substantially longer retrospective search windows than most matched local organizations despite comparatively low request frequency. The concern exists because longer lookback windows can expand the temporal scope exposed by each request even when the number of requests is modest. The leading source period contributes 31.44% and the top three contribute 60.63% of usable delivered device-hours (2025-11-01, 2025-09-01, 2025-10-01). Device-hour totals cover 79.2% of rows and are therefore observed-coverage lower bounds; alternative explanations remain testable in the review tables.

Source note: negative durations are excluded; unavailable reach values are gaps, not zero. Search time is not treated as a window-end boundary because source rows commonly use later rounded endpoints. Temporal exceptions remain in the observed totals and are separately quantified for sensitivity review.

Ventura County Sheriff

Figure 9 is the core search-scope profile for the Ventura County Sheriff label. It includes 129,109 delivered rows and 128,021 exact-distinct rows across 27 active months. The median window is 168 hours. The 90th percentile is 2,019 hours, the 99th percentile is 8,762 hours, and the maximum is 201,888 hours, about 23 years.


There are 40,815 rows at or above 720 hours and 12,476 at or above 2,160 hours. Exact-distinct requested scope totals about 687.2 billion device-hours among rows with usable device values. Device coverage is 82.88 percent, so the total is a partial-coverage lower bound rather than a complete system total. Thirteen temporal-exception rows are retained in the observed totals and separately tested; excluding them changes delivered device-hours by about 375.46 million.

The matched local comparison places the Sheriff label at the top of the displayed local cohort on frequency and high on long-window measures. That comparison is useful for prioritization, but the cohort is small and does not prove functional equivalence.

FIGURE 9

Ventura County Sheriff: Search Frequency, Scope, and Window Length

129,109 attributed rows across 27 active months • P90 window 2,019 hours • observed exact-distinct scope 687,219,385,310.13 device-hours (82.9% device-row coverage)



NCRIC

Figure 10 profiles the NCRIC label using 83,555 delivered rows and 82,595 exact-distinct rows across 27 active months. Its median window is 168 hours, P90 is 1,825 hours, P99 is 8,786 hours, and the maximum is 184,080.5 hours. Device-row coverage is 92.69 percent. One temporal-exception row changes delivered device-hours by about 12.37 million when excluded.


The most important pattern is a sustained break after mid-2024. Median monthly rows fell from 7,421 in January-July 2024 to 1,181 in August 2024-February 2026, while median observed device-hours fell about 80 percent and median networks rose from 214 to 755. This opposing movement may indicate a change in query construction, operations, export coverage, or reporting structure. Source naming remains broadly similar, but sheet-pattern differences make a source or schema explanation plausible.

The final comparison uses a 17-label activity band. It is volume-matched, not function-equivalent. Sparse intelligence-center labels are retained only as qualitative references because their row counts are too small for stable percentile claims.

FIGURE 10

NCRIC: Search Frequency, Scope, and Window Length

83,555 attributed rows across 27 active months • P90 window 1,825 hours • observed exact-distinct scope 456,092,210,422.44 device-hours (92.7% device-row coverage)



NCRIC's monthly profile shows a sustained break after mid-2024; functionally similar labels were too sparse for stable percentile benchmarking. A sustained mid-2024 change reduced attributed monthly frequency and requested device-hour scope while the typical network reach per row increased. The concern exists because the opposing movement in frequency and network reach may indicate a material change in query construction, operating practice, or source-report coverage. No source-period scope-concentration threshold was crossed. Device-hour totals cover 92.7% of rows and are therefore observed-coverage lower bounds; alternative explanations remain testable in the review tables.

Source note: negative durations are excluded; unavailable reach values are gaps, not zero. Search time is not treated as a window-end boundary because source rows commonly use later rounded endpoints. Temporal exceptions remain in the observed totals and are separately quantified for sensitivity review.

Part Five Population scale extreme windows and geographic context

The final creation responds directly to a common public question: how large is the recorded search activity relative to the population, and are the longest windows randomly scattered or concentrated in particular months? The analysis uses the 2024 American Community Survey five-year county population of 837,469 as a broad conservative denominator.

VCSO population and outlier intensity

Figure 11 reports 15,416.57 rows per 100,000 county residents over the full 27-month period, or an average of 570.98 rows per 100,000 per active month. Those are audit-log rows, not unique people. Repeated requests, broad device scope, and multiple searches related to the same matter prevent interpretation as the percentage of residents searched.

The robust-extreme threshold is about 906.34 hours, calculated as the median plus three scaled median absolute deviations. Overall, 15.28 percent of usable rows exceed that threshold. Six months in 2024 reached at least 1.25 times the full-period rate: January, February, May, July, September, and October. January had a 23.963 percent extreme rate and an observed-to-expected ratio of 1.568; September had a 23.839 percent rate and a ratio of 1.560.


The DeFlock map is included only as public geographic context. It shows displayed camera locations clustered along populated corridors and major roads. It does not establish camera ownership, VCSO access, historical availability, or the footprint of any individual request. Land-area rates are withheld because the audit rows do not identify where a query was aimed.

FIGURE 11

Ventura County Sheriff: Population Scale and Extreme Search Windows


129,109 attributed rows across 27 months • 15,417 rows per 100,000 county residents • conservative whole-county denominator

A. Recorded-search magnitude per population




These are audit-log rows—not unique people. Repeated requests and searches spanning many devices prevent interpretation as a share of residents searched.

B. Search-window magnitude on a logarithmic scale



C. Months with elevated robust-extreme rates



D. Publicly mapped ALPR camera context



Context only: displayed camera locations are not VCSO ownership, access, historical availability, or the geographic footprint of individual searches.

Source: maps.deflock.org, accessed Sept. 20, 2026; basemap © OpenStreetMap contributors.

Finding: VCSO generated a large volume of recorded searches relative to the county population, and extreme-window rates were temporally clustered rather than evenly distributed. The strongest elevated months occurred in early 2024. These patterns warrant record-level and policy review, but they do not establish who was searched, where a search occurred, or whether any individual search lacked a legitimate purpose.

Meaningful public oversight requires complete row-level audit logs; any legally required redactions should be targeted to specific protected fields rather than applied as mass or blanket redactions that prevent testing these patterns.

Method note: population uses the controlled 2024 ACS five-year county estimate. Robust extremes use median + 3 scaled median absolute deviations. The resulting 906-hour threshold is a statistical tail classification, not a finding of improper use. County population provides broad context; land-normalized rates remain withheld because audit rows do not establish search geography.

Map context: maps.deflock.org, accessed September 20, 2026; OpenStreetMap basemap.

The case for action

The evidence does not justify complacency. It shows a powerful surveillance capability operating at substantial scale while essential purpose, user, result, and geographic details remain unavailable for public testing. A system with this reach should earn continued public authorization through transparent evidence of necessity, effectiveness, proportionality, and control. The current record does not yet provide that assurance.

What the evidence directly shows

- The delivered system exports contain millions of rows and require provenance, duplicate, source-period, and partial-month controls.
- VCSO's attributed row volume is large in both absolute terms and within several comparison frames.
- The distribution of requested windows has a pronounced long tail for VCSO, VCDA, and NCRIC labels.
- For VCSO, long-window extremes are temporally clustered, with the strongest elevated rates concentrated in selected 2024 months.
- For NCRIC, a mid-2024 break changed the relationship among monthly frequency, requested device-hours, and network reach.
- Device-based scope totals are incomplete because device values are unavailable for a share of rows; reported totals are lower bounds for covered rows.

What the evidence suggests should be reviewed

- Whether very long windows were deliberately selected, produced by defaults, inherited from templates, affected by timestamp problems, or created by export logic.
- Whether the six elevated VCSO months share users, purposes, cases, organizational units, boundary values, source files, or other record-level characteristics.
- Whether mid-2024 and early-2025 changes reflect operations, policy, staffing, network configuration, vendor updates, schema changes, or shifts in which records were included.
- Whether recurring system audits, training, access reviews, sharing reviews, and retention checks were performed even when formal responsive records were not identified.

What the evidence does not establish

- The identity of a searched person or vehicle.
- The location or inhabited land area covered by a particular request.
- The number of records returned, viewed, downloaded, retained, or used in an investigation.
- The purpose, case number, or legitimacy of an individual search.
- Illegal conduct, improper surveillance, overreach, abuse, or misconduct.

Alternative explanations that remain testable

The observed patterns could reflect legitimate investigative needs, organized-retail-theft activity, changes in staffing or query practice, broader network participation, a vendor or export update, altered device-field coverage, repeated or duplicate delivery, saved-search templates, data-entry conventions, or timestamp defects. These alternatives must be tested, not assumed. If complete records confirm a lawful and necessary explanation, officials can evaluate that evidence. If the explanations cannot be documented, the public should not be asked to bear the risk of continued deployment.

Questions officials must answer

1. What written rule determines an acceptable search-window length, and when is supervisory approval required for unusually long windows?
2. What do the 30-day, 90-day, one-year, and multi-year duration clusters represent in actual platform use?
3. Can VCSO produce a field-preserving audit export in which protected narrative information is redacted narrowly while timestamps, user pseudonyms, organizational unit, purpose category, case linkage, and result counts remain testable?
4. Who performs recurring ALPR audits, how often are they performed, what sample or exception criteria are used, and where are the results documented?
5. How are user authorization and training verified if formal ALPR training records were not identified?
6. How are sharing settings, network access, device inventories, retention settings, and deletion events documented when the response says some of these are not maintained in a retrievable record format?
7. What explains the six elevated 2024 months and the mid-2024 NCRIC break?
8. What local governing body receives regular reports on use volume, long-window exceptions, audit findings, data errors, sharing, retention, and complaints?
9. Which cameras or network privileges would be removed first if the agency cannot demonstrate a specific necessity, measurable benefit, and auditable control for their continued use?

Recommended next records

Priority	Record or clarification	Public value
1	Row-level audit export with stable pseudonymous user ID, organizational unit, timestamps, purpose category, case linkage, query type, scope fields, result count, and export provenance	Allows pattern testing without requiring public disclosure of protected investigative narrative.
2	Vendor data dictionary, schema history, timezone rules, and export-change log	Separates operational changes from reporting artifacts.
3	Long-window policy, defaults, saved-search templates, supervisory approval rules, and exception reports	Explains whether the tail reflects intended practice, defaults, or errors.
4	Recurring audit reports, access reviews, training rosters, and corrective-action records	Tests whether written safeguards were implemented.
5	Current network-sharing configuration, change history, and agency-access roster	Clarifies which data sources were available without inferring use from label names.
6	Device inventory with owner, location, active dates, funding source, and contract-city attribution	Supports deployment accountability while remaining separate from query geography.
7	Retention configuration, deletion logs, legal-hold rules, and reconciliation of local policy with vendor defaults	Allows the public to distinguish image retention from archived evidence and audit-log retention.

Recommended redaction principle

Lawful confidentiality and public auditability do not have to be treated as all-or-nothing choices. If a narrative reason, plate, case detail, or personal identifier is protected, the agency can redact that field while preserving non-sensitive structure such as timestamps, stable pseudonyms, unit codes, purpose categories, duration, scope, result counts, and source provenance. Mass or blanket redaction removes the very fields needed to determine whether concerning patterns have a routine explanation.

Decision standard for continued use

Continued deployment should not be the automatic default. Each camera, data-sharing relationship, and network privilege should be supported by a documented public-safety purpose, a measurable benefit, a defined geographic and operational need, enforceable limits, complete auditability, and periodic public review. Where those conditions are not met, the appropriate action is removal, disconnection, or non-renewal. Where an immediate decision is not possible, expansion should remain paused while the evidence is completed.

Technical methods

Population and observation unit

The analytical unit is a delivered audit-log row attributed to an organization label. Delivered rows preserve the source record population. Exact-distinct rows remove only records that match the defined fingerprint exactly. Neither measure represents unique people, unique vehicles, unique investigations, or unique users unless a separate field explicitly supports that interpretation.

Window duration

Window duration is calculated from the reported query-window endpoints. Negative durations are excluded from duration statistics. Search time is not treated as a hard window-end boundary because reviewed source rows commonly use later rounded endpoints. Records whose window endpoints extend beyond a reporting or claimed source period are retained in observed totals and reported as temporal exceptions for sensitivity review.

Requested device-hours

Requested device-hours multiply reported window length by reported devices searched for rows with usable device values. This is a scope measure. It is not a count of detections, records returned, people, vehicles, or retained images. Rows with absent or raw-missing device values are unavailable and are not treated as zero. Coverage percentages accompany the totals.

Robust extremes

For Figure 11, the extreme threshold is the median duration plus three scaled median absolute deviations. With a 168-hour median and a 166-hour median absolute deviation, the controlled threshold is approximately 906.34 hours. The method is resistant to the influence of the largest values, but the label “extreme” remains statistical. It does not establish improper use.

Exposure-adjusted monthly comparison

Monthly extreme rates divide extreme rows by usable rows for the same month. Expected counts use the full-period extreme rate multiplied by each month’s usable row count. The observed-to-expected ratio therefore controls for monthly row volume. The displayed elevated-month rule is at least 1.25 times the full-period rate.

Peer comparisons

Peer frames include county sheriffs, a volume-matched activity band, a population band, and a nationwide sheriff universe. Percentiles are tie-aware and descriptive. Organization type and geography classifications remain subject to human reconciliation. Volume-matched groups are not described as function-equivalent.

Geographic context

Publicly mapped camera markers are used only to show visible corridor context. The map does not establish camera ownership, agency access, the historical camera population available during a query, or the physical footprint searched. Total land area and inhabited land are not used as denominators because the logs do not identify query geography.

Validation status

- Database modified by supplemental scripts: no.
- Raw evidence modified: no.
- Classification register modified by the final visual generators: no.
- Publication report automatically modified: no.
- Dashboard automatically modified: no.
- Public release authorized by the analytical checkpoint: no; human and external review remain required.

Analysis output index

The investigation produced figures and review tables rather than a single opaque score. The following index explains the role of each output family and the question it is intended to answer.

Output family	Primary question	Key caution
Expanded monthly trends and source coverage	How did volume, labels, network breadth, and window behavior change over time?	Partial months and export-population changes can resemble operational changes.
Geography and regional profiles	Which approved or candidate labels are associated with states or regional entities?	A label does not prove search direction, sharing, or jurisdiction.
Non-California profile and device-hour concentration	Which labels combine high frequency, long windows, persistence, or broad requested scope?	Candidate classification and incomplete device coverage require qualification.
VCSO peer suite	Does VCSO's relative position persist across reasonable comparison frames and weightings?	Peer membership and weighting affect percentile and composite values.
Ventura-area label profile	How do similarly named local labels differ?	Name similarity is not organizational identity.
VCDA, VCSO, and NCRIC scope profiles	How are frequency, duration, reported reach, and requested device-hours distributed for each label?	Scope is not returned results, affected persons, or legality.
NCRIC regime and peer audit	Does the mid-2024 break coincide with source, schema, or comparison-cohort concerns?	The audit identifies plausible explanations but does not establish cause.
VCSO intensity preflight and Figure 11	How large is row activity relative to population, and are statistical extremes temporally clustered?	Population is a broad denominator; row rates are not resident-search rates.

Named analytical files generated for follow-up

- Monthly search-scope profiles and duration distributions.
- Source-provenance and label-reconciliation tables.
- Extreme-window row tables and temporal-exception tables.
- Value-state profiles distinguishing numeric, absent, and raw-missing fields.
- Peer-membership and function-reference tables.
- NCRIC regime-source audit and change-boundary rows.
- VCSO population denominator review, robust outlier profile, extreme rows, and monthly extreme concentration.

Sources reviewed

- Ventura County Sheriff's Office response to public-records request, dated September 14, 2026.
- VCSO Policy 429, Automated License Plate Readers, dated December 23, 2022.
- Flock Safety contract and related order, grant, and implementation materials supplied with the response.
- BSCC Organized Retail Theft Prevention Grant budget workbook supplied with the response.
- Monthly Organization Audit and Network Audit exports inducted into the controlled analytical project.
- 2024 American Community Survey five-year county population estimate used in the controlled denominator review.
- Maps.deflock.org geographic context, accessed September 20, 2026, using an OpenStreetMap basemap.

Call to action

The public should not have to accept an expansive surveillance system on the strength of incomplete exports, unavailable configuration records, undocumented training, and assurances that cannot be independently tested. The present analysis identifies concrete periods, duration patterns, data-quality exceptions, and governance gaps that officials can investigate now.

Praeva Analytics calls for an independent investigation, a pause on expansion and renewal commitments, targeted production of complete audit records, and a public necessity review of every camera and network privilege. The end point of that review should be clear: retain only what can be proven necessary, effective, proportionate, and auditable; remove or disconnect what cannot.