

LAKANA SOS

SOVEREIGN OPERATING SYSTEM FOR CIVILIAN SAFETY

Privacy-Preserving Civilian Safety Infrastructure for Protection Without Surveillance

A Local-First, Fail-Closed Architecture for Safety, Evidence Custody, and Resilient Communication Without Continuous Behavioral Monitoring

MarTaize K. Fails
Founder, LAKANA Sovereign Systems LLC
Norman, Oklahoma, USA

September 2026
Public White Paper / SSRN Working Paper v2.1

Core thesis. *Civilian protection should not require persistent third-party observability of ordinary life. LAKANA SOS - the LAKANA Sovereign Operating System for civilian safety - tests whether protective capability can be preserved while materially reducing continuous behavioral exposure.*

Companion technical preprint. This white paper is the public-facing companion to the LAKANA SOS v71 technical preprint archived on Zenodo, DOI 10.5281/zenodo.19956214. The technical preprint is the controlling source for simulation methods, statistical definitions, exact quantitative results, exclusions, evidence tiers, and reproducibility details. This white paper interprets the problem, architecture, governance doctrine, and evidence boundaries at a public-safe, non-reconstructive level.

Scope statement. This document is not a deployment certification, emergency-response certification, medical-device claim, law-enforcement suitability claim, independent privacy/security certification, or guarantee of real-world survival. It is a public architecture and validation-readiness white paper supported by bounded simulation evidence and companion technical reporting.

Revision note (v2.1). This release standardizes the public version designation and incorporates minor editorial clarifications. Quantitative results, evidence tiers, limitations, and claim boundaries are unchanged.

Abstract

Civilian safety technologies increasingly rely on continuous location collection, cloud telemetry, networked cameras, centralized search, and behavioral analytics. These capabilities can support legitimate protective purposes. The unresolved design question is whether obtaining that value should require making ordinary human activity persistently observable to external systems. Persistent observability is not merely a privacy inconvenience: it creates a durable information asset that can reveal routines, relationships, sensitive locations, vulnerabilities, and deviations from ordinary behavior while expanding the attack surface for misuse, compromise, legal compulsion, secondary use, and future inference.

LAKANA SOS - the LAKANA Sovereign Operating System for civilian safety - proposes a different design objective. Critical safety state should remain local where technically possible; raw observations should be reduced toward bounded safety-relevant state rather than automatically becoming persistent remote telemetry; uncertainty should constrain authority rather than justify broader inference; evidence should be constructed around qualified safety events; and disclosure should be purpose-bound, consent-aware, time-limited, auditable, and revocable where the context permits. The objective is not to eliminate sensing, evidence, communication, or institutional response. It is to reduce the amount, persistence, centralization, queryability, and secondary-use potential of information required to provide those functions.

The paper examines family-location systems, connected-vehicle telemetry, location brokerage, home-security access, tracking misuse, automated license plate reader networks, and behavioral profiling as structural case studies rather than indictments of particular companies. Oklahoma provides a live governance contrast: Norman declined municipal Flock adoption despite police interest in its investigative and safety value, while Oklahoma City retained a large deployment with additional restrictions. The point is not to declare either city or vendor correct; it is to show that communities still lack a settled architecture for retaining useful safety capability while limiting persistent observability.

The companion LAKANA SOS v71 technical preprint reports architecture-level Monte Carlo evidence under declared degraded and adversarial assumptions. Those results do not establish field effectiveness, certification, legal compliance, or universal superiority. This white paper advances a narrower, falsifiable proposition: civilian safety architectures should be evaluated simultaneously for protective capability and surveillance burden, and where comparable declared protective requirements can be met with materially less persistent behavioral exposure, the lower-burden architecture should generally be preferred.

Keywords: civilian safety; privacy-preserving infrastructure; Sovereign Operating System; safety without surveillance; local-first computing; persistent observability; location privacy; automated license plate readers (ALPR); behavioral profiling; evidence custody; degraded communications

How to read the publication pair

Document	Primary burden
This white paper	Why persistent observability creates safety and governance risk; public-safe LAKANA architecture; consent/evidence doctrine; bounded interpretation of the technical evidence; validation path.
LAKANA SOS v71 technical preprint	Threat/failure model; empirical and proxy inputs; Monte Carlo design; statistical methods; TSARO/NICOLE diagnostics; exclusions; reproducibility and evidence-tier details.
Controlling rule	Both documents use the same evidence boundaries. The v71 technical preprint controls quantitative definitions, methods, assumptions, and exact simulation results.

Contents

1. Executive Summary
 2. The Design Problem: Safety and Surveillance Have Been Unnecessarily Coupled
 3. From Safety Data to Surveillance Infrastructure: The Structural Risk Landscape
 4. Why Persistent Observability Is a Safety Risk
 5. A Formal Design Objective: Surveillance Burden and the Safety-Surveillance Frontier
 6. What LAKANA SOS Is
 7. Architectural Method: How LAKANA Reduces Surveillance Dependencies
 8. Blue Force Bridge: Consent-Gated Rescue Communication
 9. Evidence Base: What the v71 DOI Supports
 10. Evidence-Tier Discipline and Claim Boundaries
 11. Falsifiable Research Questions and Validation Metrics
 12. Initial Pilot Domains
 13. Governance, Anti-Aggregation, and Non-Weaponization Requirements
 14. Legal and Institutional Considerations
 15. Limitations
 16. Validation and Translation Roadmap
 17. Conclusion: Protection Without Persistent Observability
- Statements and Declarations
- References

1. Executive Summary

The civilian safety market has often treated continuous visibility as the price of protection. Family-tracking, crash-detection, lone-worker, campus-safety, home-security, connected-vehicle, and public-safety systems can begin with legitimate protective purposes and still generate records whose persistence, searchability, and reuse extend far beyond the immediate event.

LAKANA begins from a different premise: privacy can itself be a safety property. A location history can expose a domestic-violence survivor's escape route, a child's custody exchange, a journalist's source meeting, a patient's clinic visit, a worker's organizing activity, or the recurring pattern that tells an attacker when a home is empty. The relevant question is therefore not only whether information is protected after collection, but whether the protective function required creating the reusable record at all.

LAKANA SOS - the LAKANA Sovereign Operating System for civilian safety - is a local-first, fail-closed architecture intended to preserve bounded safety state, evidence integrity, degraded-communications resilience, user authority, and consent-gated institutional integration without turning ordinary activity into a generalized behavioral dossier. TSARO constrains action under uncertainty, NICOLE governs incident evidence, CivOS preserves local execution and state, and the transport layer attempts permitted alternatives when normal infrastructure fails.

The public-interest argument is intentionally vendor-neutral. Networked ALPRs, family-location tools, connected vehicles, data brokers, home-security systems, and behavioral analytics are discussed because they expose recurring design questions: when does useful observation become historical reconstruction, when does live visibility become interception capability, and when does safety data become a profile that affects unrelated decisions? Norman and Oklahoma City illustrate that these questions remain unresolved even when public-safety officials identify real operational value.

The companion v71 technical preprint carries the quantitative burden: model assumptions, Monte Carlo design, statistical intervals, empirical/proxy evidence tiers, diagnostics, exclusions, and reproducibility. This white paper carries a different burden: explain the problem, show how LAKANA is architecturally different, state the governance non-negotiables, interpret the evidence without overclaiming, and define the external tests that could validate or falsify the protection-without-surveillance proposition.

Persistent-observability pattern



LAKANA SOS — Sovereign Operating System design objective



Design principle Minimize creation of reusable behavioral records while preserving event-level protective capability.

Collection minimization is evaluated before access control: the first question is what ordinary-life data must exist at all.

Figure 1. Conceptual contrast between a persistent-observability architecture and LAKANA SOS's event-bounded design objective.

2. The Design Problem: Safety and Surveillance Have Been Unnecessarily Coupled

The architecture of connected safety developed for understandable reasons. Cloud computing made centralized processing inexpensive. Smartphones made continuous location available. Networked cameras made physical observations searchable. Machine learning made inference increasingly scalable. Large databases made historical

reconstruction practical. Each development created real capability. Taken together, however, they encouraged a default assumption that **visibility and safety should rise together**.

To protect a family member, know where the family member is. To protect a driver, collect how the driver moves. To locate a stolen vehicle, observe vehicles moving through a camera network. To investigate an event, retain historical observations. To detect danger, continuously send enough information to a remote system that the remote system can identify it. The pattern is coherent, but it is not technically inevitable.

A safety architecture can instead distinguish among **ordinary life**, **uncertain safety conditions**, and **qualified safety events**. During ordinary life, information should remain maximally local and minimally persistent. During uncertain conditions, local evidence may accumulate without automatically widening disclosure. During a qualified event, narrowly bounded evidence may be released for a defined protective purpose. This state separation changes the engineering question from 'How much information can we collect so that we are prepared for any future event?' to 'What is the minimum information and authority required at this state to preserve safety?'

This paper therefore treats surveillance not only as a civil-liberties concern but as a system property. A design has surveillance consequences when its routine operation makes ordinary people externally observable, historically reconstructable, linkable across contexts, searchable by additional actors, or classifiable by models in ways that are not necessary for the immediate protective function. That property can exist even when every operator is well intentioned.

As a design principle, this paper places the burden of justification on persistent observability. When comparable protective capability can be achieved with less collection, shorter retention, less centralized custody, weaker historical queryability, fewer secondary uses, and stronger user authority, the lower-surveillance architecture should generally be preferred.

3. From Safety Data to Surveillance Infrastructure: The Structural Risk Landscape

3.1 The problem is structural, not vendor-specific

The examples in this section are not presented to establish that a named company, government, police department, insurer, application developer, or technology category is inherently malicious. Many systems discussed here were created to solve real problems. ALPRs can locate stolen vehicles and vehicles associated with missing persons. Family-location applications can help caregivers find relatives. Telematics can encourage safer driving. Home-security cameras can preserve evidence after violent incidents. Connected-car services can call emergency responders after crashes. These benefits matter.

The relevant inquiry is different: **what additional power is created when safety-relevant observations become persistent, centralized, queryable, linkable, and reusable?** An architecture can serve a legitimate purpose and still create a secondary surveillance capability. That capability may emerge without any initial intent to misuse the system. The correct analysis therefore starts before the privacy policy, access rule, or disciplinary process and asks why the system possessed enough information and authority for a later misuse or repurposing to be possible.

3.2 Continuous collection and the creation of latent surveillance assets

Data does not need to be actively abused to have surveillance significance. Persistent datasets create **latent capability**. They may later become searchable with better algorithms, linkable to new databases, accessible under new legal authorities, transferable through acquisition, valuable for new commercial purposes, or vulnerable to credential compromise. Privacy risk therefore cannot be evaluated only by asking what an operator does with information today. Architecture must also consider what the existence of the information makes possible tomorrow.

Family-location systems illustrate the tension. Investigative reporting described Life360's historical sale of precise location data to multiple data brokers; the company later announced changes that reduced certain precise-data sales [1-2]. The point is not to freeze a company in an older business model. The point is that data generated in the name of family safety can become economically valuable precisely because it is persistent, precise, and behaviorally rich.

Driver telemetry demonstrates the same migration risk. The Texas Attorney General sued Allstate and Arity in 2025, alleging collection and use of mobile location and driving behavior data through third-party applications for insurance-

related purposes; the defendants disputed wrongdoing [3-4]. In 2026 the FTC finalized an order settling allegations that GM and OnStar collected and sold precise geolocation and driving behavior data without adequate informed consent, including disclosures to consumer-reporting agencies [14]. These matters show why safety-related telemetry should not silently become a generalized scoring asset.

3.3 Sensitive-location brokerage and the inference problem

FTC enforcement against X-Mode/Outlogic, Gravy Analytics/Venntel, and Mobilewalla demonstrates that precise location can expose substantially more than a coordinate. The Commission described risks involving medical and reproductive-health facilities, places of worship, domestic-abuse shelters, labor union offices, political gatherings, military installations, and other sensitive places [5-7]. Gravy/Venntel was also alleged to derive sensitive characteristics including medical decisions, political activity, and religious viewpoints from location information [6].

The consequence is architectural: **location is an inference channel**. A safety system that retains generalized movement data may unintentionally create information about health, belief, association, family status, work, political activity, or vulnerability without ever asking the person to disclose those attributes directly.

3.4 Home-security, responder access, and internal authority

The FTC's 2023 Ring matter shows a related risk: safety and security devices can create privacy failure through overly broad internal access and inadequate controls. The Commission alleged that employees and contractors had excessive access to private videos and that security failures enabled account and camera compromise [8]. The relevance to LAKANA is not that responder access is inherently unsafe. It is that access authority must be **rescue-scoped**, time-bounded, logged, and separated from ordinary monitoring.

The same principle applies to any future Blue Force Bridge implementation. A responder may need location, audio, status, or scene context during a legitimate emergency. That does not justify permanent live visibility or unrestricted institutional access during ordinary life.

3.5 Networked ALPR infrastructure as a contemporary architectural case

Automated license plate recognition is particularly instructive because it demonstrates the difference between an isolated public observation and a network capable of retrospective reconstruction. One camera recording one vehicle at one location is limited. A dense network operating across time, jurisdictions, and searchable systems creates a materially different capability: an authorized user can ask where a vehicle has appeared, identify recurring routes or locations, and potentially combine that information with other records.

Flock Safety is a prominent contemporary example, but this paper's critique is directed at the architecture category, not the company. Flock states that approved users must authenticate, searches require defined investigative purposes, queries are logged, supervisors can review activity, network sharing is customer-controlled, and customers determine retention [25]. Flock also describes seven-day default retention for new law-enforcement deployments, evidence-preservation mechanisms for active cases, customer-data restrictions, and policy limits on license-plate-reader use [26-27, 40]. These are material safeguards and should be acknowledged.

Those safeguards govern an already-created data asset. Authentication governs who can access it. Audit logs govern whether access can later be reviewed. Retention limits govern how long it persists. Sharing controls govern which institutions may query it. None of those controls by themselves answer the earlier question: **was persistent collection of non-event movement necessary to deliver the protective function?**

Flock's own 2026 disclosure regarding California illustrates why architecture and governance must be examined together. The company reported that some California agencies' camera networks had been inadvertently accessible to out-of-state law-enforcement agencies during 2025 and stated that earlier limitations in technical logging made the precise cause impossible to determine in some instances; Flock subsequently implemented additional guardrails [26]. The appropriate conclusion is not that safeguards are meaningless. It is the opposite: complex sharing systems require safeguards because the underlying data has substantial power once collected.

3.6 Oklahoma as a live governance experiment: Oklahoma City and Norman

Oklahoma provides an unusually instructive local comparison because neighboring jurisdictions confronted similar public-safety questions and reached different conclusions.

In Norman, the disagreement was **not** over whether ALPR technology could help police. Mayor Stephen Tyler Holman told 404 Media that Flock made two presentations over three years at the request of the Norman Police Department, which wanted to implement the system. The police case emphasized crime prevention, more efficient investigations, Amber Alerts, Silver Alerts, kidnapped children, and locating vulnerable older adults. Norman's elected leadership nevertheless rejected municipal adoption; all nine council members opposed the most recent pitch after concerns about retention, access, security, private custody, and collection of records on vehicles whose drivers had committed no crime [28].

Norman therefore confronted substantially the same research question examined in this paper: **can a community give police and responders effective tools and create a safer environment without making persistent surveillance of ordinary residents the architectural price of that assistance?** The fact that the technology was useful to police makes the question harder and more important, not less important. If a technology had no safety value, rejecting it would require little analysis. The meaningful engineering challenge is to preserve the useful function while materially reducing the persistent observation that accompanies it.

Oklahoma City reached a different policy judgment. On August 18, 2026, the City Council voted 5-3 to renew a \$270,000 contract covering 90 Flock cameras [29]. The police department described additional safeguards, including shorter seven-day retention and reduced national sharing; local reporting also noted that 109 local agencies could access Oklahoma City data under then-current sharing arrangements [30]. The contrast demonstrates that communities can share the same public-safety objective while disagreeing about the acceptable surveillance burden.

Norman's municipal rejection also demonstrates the limits of jurisdiction-specific privacy policy. Cleveland County and private homeowners associations may operate separate Flock systems in or around Norman even though Norman city government does not [28]. Privacy can therefore be weakened by neighboring jurisdictions, private infrastructure, interoperability, or later governance changes. LAKANA treats sovereignty as a systems property rather than a promise made by one institution at one point in time.

The policy debate is broader than Oklahoma. In September 2026, Florida barred local law-enforcement agencies from installing license-plate readers on state highways, reflecting wider reevaluation of how network scale, retention, and cross-jurisdiction access should be governed [31]. This paper treats such policy changes as evidence of an unsettled design problem, not as proof that any one jurisdiction has resolved it.

Risk class	Legitimate safety value	Structural risk	LAKANA design implication
Family location	Caregiver awareness; missing-person support	Routine and family movement becomes a persistent asset	Keep ordinary state local; event-bound release
Vehicle / telematics	Crash response; driving feedback	Behavior can migrate into insurance or economic scoring	Separate safety state from generalized driving profile
Location brokerage	Analytics and operational intelligence	Sensitive places and associations can be inferred	Minimize non-event egress and cross-context linkability
Home / responder access	Evidence and emergency context	Broad internal access can become continuous visibility	Rescue-scoped, logged, time-bounded access
Networked ALPR	Stolen vehicles; missing persons; investigations	Ambient collection enables retrospective movement reconstruction	Evaluate protective function against historical queryability
Behavioral profiling	Risk estimation; personalization	Profiles can influence consequential decisions and may be wrong	No hidden safety-to-scoring purpose migration; fail closed under uncertainty

4. Why Persistent Observability Is a Safety Risk

Searchable location creates retrospective power: "Where were you?" Persistent surveillance creates contemporaneous power: "Where are you now?" Behavioral profiling creates inferential power: "What does your behavior imply about you?" The three powers overlap, but each creates distinct safety and governance risks.

4.1 Privacy loss is not merely an informational inconvenience

Debates about surveillance are often framed as a trade between an abstract privacy interest and a concrete safety interest. That framing is incomplete. Loss of privacy can itself create physical, economic, psychological, associational, and institutional safety risks.

A person's location history may reveal where the person sleeps, works, receives medical treatment, worships, attends political events, exchanges custody of children, seeks counseling, or spends nights away from home. Repetition enables additional inferences: which address is probably home, when it is usually vacant, who repeatedly appears at the same places, which routes are habitual, and what constitutes an unusual deviation from routine.

The danger does not begin only when a database is demonstrably abused. It begins when the architecture creates a durable capability to reconstruct ordinary life. Once that capability exists, legitimate users may misuse access, credentials may be compromised, policies or legal authorities may change, data may migrate through acquisition or secondary use, and future analytical systems may extract conclusions that were impossible when the information was first stored.

LAKANA therefore distinguishes protection of information after collection from minimization before information becomes a surveillance asset. Encryption, authentication, retention limits, audit logging, and access controls remain important, but they govern an existing record. A perfectly encrypted record can still reveal a person's life to an authorized reader; a perfectly audited query can still be unnecessary; a seven-day movement history can still expose a sensitive seven days.

4.2 Searchable location creates retrospective power

Searchable location data changes the nature of observation. Seeing a vehicle at one intersection is an observation. Being able to enter an identifier and reconstruct where that vehicle has appeared over a meaningful period is a different capability. The distinction is between witnessing an event as it happens and reconstructing a person's past after deciding that the person has become interesting.

This retrospective capability removes an important historical constraint on surveillance: investigators traditionally needed some reason to watch a specific person before the relevant conduct occurred. Large searchable repositories can invert that sequence. Collection occurs first. Interest may arise later. The stored history can then be queried retroactively.

The Supreme Court confronted a related transformation in **Carpenter v. United States**, 585 U.S. 296 (2018). The government obtained 12,898 historical cell-site location points covering 127 days of Timothy Carpenter's movements. The Court held that acquisition of the historical CSLI was a Fourth Amendment search and emphasized the revealing nature of a comprehensive digital chronicle of physical movement [19]. Carpenter did not involve ALPRs and does not automatically determine the constitutionality of camera networks. Its systems insight is directly relevant: aggregation changes the informational character of individually ordinary location records.

In **United States v. Jones**, 565 U.S. 400 (2012), the Supreme Court held that government installation and use of a GPS tracking device on a vehicle constituted a search. The majority relied on physical trespass, while concurring opinions addressed the broader privacy implications of long-term electronic monitoring [20]. Together, Jones and Carpenter reflect judicial recognition that technologically inexpensive, longitudinal tracking can create a qualitatively different form of governmental knowledge than episodic visual observation.

4.3 ALPR scale and the mosaic problem

The Massachusetts Supreme Judicial Court addressed ALPRs directly in **Commonwealth v. McCarthy**, 484 Mass. 493 (2020). The particular network - four fixed readers at two Cape Cod bridges - was too limited on the record before the court to constitute a constitutional search. But the court refused to establish a categorical rule and recognized that

sufficiently widespread ALPR deployment could implicate constitutional privacy by revealing a substantially fuller picture of a person's public movements [22].

The significance for engineering is that **scale changes the capability**. One camera is not analytically equivalent to a dense interoperable network. The risk grows with observation density, retention, identity linkage, historical queryability, interoperability, and analytical capability.

The ongoing **Schmidt v. Norfolk** litigation makes the issue contemporary. Plaintiffs challenge Norfolk, Virginia's use of a Flock-supplied ALPR system under the Fourth Amendment; in April 2026 the ACLU, ACLU of Virginia, and EFF filed an amicus brief in the Fourth Circuit arguing that interconnected ALPR systems enable sensitive retrospective searches. The case remains ongoing, so this manuscript does not present the constitutional argument as an adjudicated final rule [24]. The unresolved litigation is itself evidence that the governance problem is not settled.

4.4 Searchable location is dangerous even when the person querying it is not violent

The most severe use case is stalking or physical violence, but the autonomy harm begins earlier. Consider a former partner who is not presently violent but is jealous, controlling, suspicious, intrusive, or simply unwilling to accept that another adult's movements are no longer their business. Searchable location can answer questions that person has no legitimate right to answer: Where did she sleep? Whose house does he visit after work? Was she actually at work? Where does a new partner live? Did he visit a lawyer? Where do the children exchange custody?

A surveillance system does not need to cause physical injury for these queries to become coercive. Repeated checking can facilitate confrontation, accusations, harassment, custody conflict, obsessive monitoring, or coercive control. And because violent behavior can emerge after a period of nonviolent monitoring, an architecture cannot safely assume that a currently trusted person will remain harmless indefinitely.

Federal prosecutions show how location technology can become part of escalating intimate-partner abuse. In the Victor Ortiz case, investigators found that Ortiz had attached a GPS device to his former girlfriend's vehicle to monitor and follow her; he was later convicted after a caustic chemical attack that permanently injured her and was sentenced to seven years in federal prison [9]. In an earlier District of Columbia prosecution, Ulus Jorden hid a GPS device inside a stuffed toy given to his daughter so he could track his former girlfriend; he was convicted of stalking [10]. These cases do not imply that ordinary location-sharing users are violent. They demonstrate that the same technical capability can become physically dangerous when control relationships change.

4.5 Insider access converts institutional trust into individual power

Sensitive systems are often defended on the ground that access is restricted to trusted professionals. That control is necessary but insufficient: every authorized account is another credential, endpoint, judgment call, and potential misuse path.

Former Irondequoit police officer William Rosica provides a documented example. Federal prosecutors established that he used law-enforcement databases and restricted systems to obtain information about a former girlfriend and her family as part of a broader stalking and harassment campaign; he pleaded guilty to cyberstalking and computer intrusion and received a five-year sentence [11]. The institutional databases had legitimate purposes. The abuse arose because a person with authorized institutional position could convert organizational capability into personal power.

Audit logs may eventually reveal misconduct, but the victim cannot become un-located after the query occurs. LAKANA therefore treats after-the-fact accountability as necessary but weaker than least-authority architecture that limits what any credential can expose.

4.6 External compromise creates the same danger without an insider

The threat model also includes people who were never authorized. In 2023 federal prosecutors charged two men with using compromised law-enforcement credentials to access a restricted federal portal and then using obtained personal information in an extortion scheme [12]. A searchable sensitive database therefore has not only its intended users, but also anyone who may steal, coerce, socially engineer, or otherwise acquire their authority.

The FTC's SpyFone case provides a consumer-surveillance analogue. The Commission alleged that the stalkerware service enabled covert access to location and other sensitive information and that inadequate security allowed unauthorized access to consumer data [13]. Collecting sensitive location information creates a cybersecurity asset that

must be defended for as long as it exists; not collecting unnecessary information removes that particular stored-data breach target.

4.7 Persistent surveillance creates contemporaneous power

Searchable historical location answers **Where was this person?** Persistent real-time or near-real-time surveillance answers **Where is this person now?** That difference materially increases physical risk because historical data supports reconstruction while live visibility can support interception.

A hostile actor who knows a person's present location does not need to infer tomorrow's routine; the actor can travel to where the person is now. This matters in stalking, domestic violence, witness intimidation, theft, harassment, organized crime, and targeting of other vulnerable people.

Real-time visibility can also expose a protective strategy. A survivor temporarily relocating, a threatened witness changing routes, a journalist meeting a source, or a person staying with a friend after harassment may depend on unpredictability. Persistent observability erodes that defensive uncertainty.

4.8 Persistent surveillance can reveal sensitive places and associations

FTC enforcement repeatedly identifies sensitive-location exposure as a concrete harm. X-Mode/Outlogic's order addressed precise location data capable of showing visits to medical and reproductive-health clinics, places of worship, and domestic-abuse shelters [5]. Gravy/Venntel and Mobilewalla matters described location-derived information involving health care, religious practice, labor union activity, political gatherings, military installations, and other sensitive contexts [6-7].

Association can itself be safety-relevant. The Supreme Court's decision in **NAACP v. Alabama ex rel. Patterson**, 357 U.S. 449 (1958), did not involve digital surveillance, but it recognized the relationship between privacy in association and freedom of association in a context where disclosure exposed members to economic reprisal, employment loss, coercion, and hostility [23]. Modern location systems can sometimes infer association without ever requesting a membership list.

The resulting harm includes a chilling effect: the support group not attended, protest skipped, source not contacted, clinic not visited, union meeting avoided, or religious service attended less freely because presence may later be reconstructed. Part of the injury is therefore lawful activity that never occurs.

4.9 Persistent aerial surveillance demonstrates how aggregation changes public observation

The Fourth Circuit's en banc decision in **Leaders of a Beautiful Struggle v. Baltimore Police Department**, 2 F.4th 330 (4th Cir. 2021), involved Baltimore's aerial surveillance program. The court held that accessing the program's retained data constituted a Fourth Amendment search because the system could deduce from the whole of individuals' movements; the court reversed the denial of an injunction against warrantless access [21].

The case is especially relevant because each individual observation occurred in public. The constitutional concern emerged from **aggregation and retrospective reconstruction**, not from converting a private room into a camera feed. This is the same systems distinction that matters for networked roadway observation.

4.10 Behavioral profiling creates inferential power

The most consequential evolution is from tracking where someone is to deciding what the data means about them. Behavioral profiling can combine location, movement, browsing, purchasing, communications, device signals, demographics, and interaction data to classify people or generate predictions that influence price, opportunity, scrutiny, targeting, or access.

This produces a third form of surveillance power: searchable location reconstructs the past; persistent surveillance observes the present; behavioral profiling attempts to classify or predict the person. The last step is especially consequential because an inferred profile can affect someone even when the inference is wrong.

4.11 Profiling can convert safety data into insurance and economic consequences

Connected vehicles provide a clear example of purpose migration. The FTC alleged that GM and OnStar collected precise geolocation and driving behavior - including hard braking, speeding, and late-night driving - and disclosed data

to consumer-reporting agencies without adequate informed consent. The final 2026 order restricts such sharing and requires additional consumer controls [14].

The lesson extends beyond automobiles. A feature perceived as navigation, crash assistance, driving feedback, or safety support can generate a behavioral profile capable of affecting an unrelated economic relationship. For LAKANA, safety information should not quietly become pricing information.

4.12 Profiling can expose health conditions and intimate states

Health information need not contain an explicit diagnosis to become health-related intelligence. Visits, searches, medications, questionnaire responses, location patterns, and repeated behavior can collectively expose sensitive conditions.

The FTC's BetterHelp matter alleged that the online counseling service shared email addresses, IP addresses, and answers to health questionnaires with advertising platforms despite privacy assurances; the final order prohibited sharing health data for advertising and required \$7.8 million for consumer refunds [15]. The broader risk is portability: once sensitive behavioral information crosses contexts, it can expose facts users never intended to disclose outside the immediate service relationship.

4.13 Profiling can reveal political, religious, union, and social identity

The FTC alleged that Mobilewalla sold audience segments based on sensitive characteristics such as medical conditions and religious beliefs and handled data associated with political gatherings, labor-union offices, and other sensitive places [7]. Gravy/Venntel was alleged to derive medical decisions, political activity, and religious viewpoints from location data [6].

This is categorically different from storing coordinates. A profiling system creates propositions - likely patient, worshipper, union participant, protest attendee, or member of a sensitive segment - that the person may never have expressly disclosed.

4.14 Profiling can become a gatekeeper to fundamental opportunities

Profiles become especially dangerous when classification is connected to consequential decisions. In **United States v. Meta Platforms, Inc.**, the Department of Justice alleged that Meta's housing-advertising system used algorithms that relied in part on characteristics protected by the Fair Housing Act when determining which users received housing advertisements. A federal court approved a settlement requiring changes to the system and ongoing oversight [17].

The relevance to civilian safety architecture is the mechanism: a personalization or prediction system can become a gatekeeper. Once a profile is treated as operationally meaningful, it can influence housing, pricing, advertising, security intervention, opportunity, or institutional scrutiny.

4.15 Profiles can be wrong and still create real harm

Inference is fallible. The FTC's Rite Aid facial-recognition matter alleged thousands of false-positive matches and described consumers being followed, searched, ordered to leave stores, publicly accused of wrongdoing, or subjected to police encounters. The Commission also alleged higher false-positive burdens in stores located in plurality-Black and Asian communities [18].

The system did not need to know the person perfectly; an incorrect classification became consequential because another actor trusted it. A profile can become operational truth before it becomes factual truth. This directly supports LAKANA's fail-closed doctrine: weak, incomplete, or contradictory evidence should surface uncertainty rather than silently expand authority.

4.16 Function creep is predictable, not exceptional

Function creep can arise without malicious intent. A system is built for Purpose A, accumulates useful data, and later discovers that the same data can serve Purpose B or C.

The FTC's Avast matter illustrates the pattern in another domain. The Commission alleged that software marketed as protecting users from online tracking sold re-identifiable browsing data through a subsidiary to more than 100 third parties; the final order restricted such sales and required a \$16.5 million payment [16]. The engineering lesson is

straightforward: data that already exists invites additional uses, so architecture should minimize unnecessary reusable data rather than rely entirely on future restraint.

4.17 Profiling can create feedback loops

Behavioral profiling can create self-reinforcing loops: a person, location, or pattern is classified as higher risk; that classification causes more monitoring; more monitoring produces more observations; and the larger observation set can appear to confirm the original classification simply because the system watched the target more intensely.

The loop is classification -> additional observation -> additional data -> apparent confirmation -> stronger classification. In safety-critical settings, such loops can influence police attention, employer decisions, insurance costs, parental monitoring, or reputation. Architecture must therefore distinguish independent evidence from artifacts created by its own observation policy.

4.18 The legal authorities reveal a common technological principle

The authorities discussed above arise under different laws, doctrinal tests, and factual circumstances. They should not be collapsed into a claim that courts have declared persistent surveillance categorically unlawful. Carpenter concerned historical CSLI. Jones involved a GPS device. Leaders of a Beautiful Struggle concerned persistent aerial surveillance. McCarthy concerned a limited ALPR deployment. NAACP v. Alabama concerned associational disclosure. Schmidt remains ongoing. FTC and DOJ matters arise under consumer-protection, civil-rights, cybersecurity, or criminal law.

Collectively, however, they support a more durable engineering observation: **digital aggregation changes the privacy and safety consequences of information.** Long-term movement histories reveal more than isolated observations. Real-time tracking can facilitate interception and stalking. Association disclosure can expose people to coercion and chill lawful activity. Behavioral information can migrate into insurance, advertising, housing, security, and other consequential contexts. Technological scale can transform an otherwise ordinary observation into a qualitatively different form of power.




Figure 2. Three distinct powers created by persistent observability: retrospective reconstruction, contemporaneous visibility, and behavioral inference.

4.19 A three-axis threat model for civilian safety infrastructure

Axis	Question enabled	Primary risk multipliers	Representative harms
Retrospective observability	Where have you been?	Observation density; retention; queryability; identity linkage; cross-dataset linkage	Routine reconstruction; association discovery; retroactive investigation; sensitive-location inference
Contemporaneous observability	Where are you now?	Update frequency; precision; coverage; automated alerts; recipient breadth	Stalking; interception; witness intimidation; exposure of escape or safety strategy
Inferential observability	What does behavior imply?	Modalities combined; longitudinal depth; model sophistication; cross-context reuse	Profiling; scoring; discrimination; economic consequences; false classifications

LAKANA design conclusion. *A civilian safety architecture should not be judged solely by how much danger it can detect. It must also be judged by how much power over ordinary people its operation creates.*

4.20 Public-safe architecture-to-risk mapping

Observability risk	Conventional capability	LAKANA architectural answer
Retrospective reconstruction	Searchable movement or event history	Ordinary state remains local where feasible; evidence is event-bounded; no generalized historical query surface is required by default.
Contemporaneous visibility	Persistent live location or standing dashboard access	Access is event-scoped, purpose-specific, separately authorized, logged, and designed to expire when the protective purpose ends.
Inferential profiling	Behavior aggregation, scoring, or predictive classification	Purpose isolation; no ordinary-life scoring pipeline; uncertainty narrows authority rather than silently expanding it.
Insider misuse	Standing institutional credentials to broad data	Least-authority access: narrow recipient scope, explicit purpose, authentication, auditability, and short-lived authorization.
Credential compromise	Reusable centralized visibility and historical data	Minimize externally available state so a compromised credential exposes less reusable information and less authority.
Function creep	Historical data retained because it may become useful later	Expiration, anti-aggregation, no behavioral resale, and no secondary-use pathway merely because data already exists.

The mapping is intentionally public-safe: it states what the architecture is designed to prevent without disclosing protected thresholds, anti-abuse logic, or reconstructive implementation details.

5. A Formal Design Objective: Surveillance Burden and the Safety-Surveillance Frontier

A privacy-preserving safety architecture needs a measurable target. This paper therefore defines **surveillance burden** as the degree to which a safety architecture renders ordinary human activity externally observable, persistent,

historically queryable, linkable across contexts, shareable across actors, or reusable beyond the immediate protective purpose.

Surveillance burden should be treated as a vector rather than a single universal score. Different deployments may weight dimensions differently, but the minimum reporting set should include non-event data egress, retention duration, raw-versus-derived storage, historical queryability, identity linkability, cross-context linkage, number and type of custodians, sharing breadth, secondary-use authority, user revocability, deletion verifiability, and inferential capability.

For a deployment-specific evaluation, an illustrative normalized scalar can be defined as $B_s = w_E E + w_R R + w_Q Q + w_L L + w_C C + w_S S + w_I I$, where E represents non-event egress, R retention, Q queryability, L linkability, C custody breadth, S sharing scope, I inferential power, and weights w are declared in advance. This is not a validated universal privacy metric; it is a transparent comparison scaffold under declared assumptions.

The displayed scalar intentionally compresses only seven illustrative dimensions. Revocability, deletion verifiability, raw-versus-derived storage, and other dimensions listed in the surveillance-burden vector remain separately reportable and are not implied to disappear merely because they are omitted from the illustrative scalar.

Deployments may also pre-register non-compensable constraints. For example, a maximum level of historical queryability, a minimum revocability requirement, or a prohibition on secondary behavioral resale need not be allowed to "cancel out" against lower burden elsewhere. This prevents the scalar from treating every privacy property as freely exchangeable.

The associated Safety-Surveillance Frontier asks whether declared protective-performance requirements can be maintained while surveillance burden falls. For a particular study, B_s can be minimized subject to predeclared constraints on threat handling, evidence integrity, transport success, false escalation, missed escalation, availability under infrastructure degradation, and human usability. Where a deployment satisfies the same declared protective-performance requirements with materially lower B_s , the lower-burden design should generally be preferred on the stated objective.

The central LAKANA hypothesis is therefore falsifiable: **a local-first, evidence-bounded safety architecture can move the safety-surveillance frontier by preserving useful protective capability while materially reducing continuous behavioral exposure.**

Safety-Surveillance Frontier: LAKANA research objective




Figure 3. CONCEPTUAL Safety-Surveillance Frontier - schematic only, not empirical data. The research objective is a leftward shift toward lower surveillance burden without unacceptable loss of declared protective performance.

6. What LAKANA SOS Is

LAKANA SOS means **LAKANA Sovereign Operating System**. In this paper, LAKANA SOS refers to the civilian-safety operating architecture that keeps critical safety logic as close as practical to the user, the device, and the physical event. It is intended for ordinary protection and for degraded settings where a conventional cloud-dependent application may lose cellular service, cloud reachability, device confidence, trustworthy context, or stable power.

The goal of the Sovereign Operating System is not to know everything about the user. Its goal is to preserve enough local, bounded, reviewable state to support protective action without turning the user's life into a generalized surveillance product. Sovereignty here is operational: the system should continue to enforce safety and evidence boundaries even when connectivity, external services, or institutional trust are degraded.

The architecture is therefore organized around bounded jobs rather than unrestricted observation: user control governs authority; TSARO constrains inference and escalation; NICOLE preserves incident evidence without normalizing continuous remote custody; CivOS maintains local state and execution; and the transport stack attempts permitted alternatives when normal communications fail. Cloud infrastructure can support validation, archival, and partner access without becoming the required authority path for every bounded local safety decision.




Figure 4. Public-safe LAKANA SOS layer model. Protected thresholds, anti-abuse logic, and reconstructive implementation details are intentionally not disclosed here.

7. Architectural Method: How LAKANA Reduces Surveillance Dependencies

7.1 Step 1: Local physical and device-context observation

LAKANA begins with local observations relevant to the bounded safety question. Depending on deployment, these may include device health, motion state, user input, environmental context, or permitted external signals. The

architecture does not require that every available sensor be enabled or that every observation leave the device. Collection itself is treated as an authority decision.

7.2 Step 2: Evidence reduction rather than raw extraction

Where feasible, raw observations are transformed into bounded state or incident-relevant evidence locally. The purpose is to reduce the amount of reconstructive information that must be transmitted or retained. Derived state is not automatically harmless, so it remains governed, but local reduction can materially shrink exposure relative to persistent raw telemetry.

7.3 Step 3: TSARO deterministic risk-envelope evaluation

TSARO is the system's deterministic skepticism and risk-envelope layer. Publicly, its important property is behavioral: missing evidence narrows action, contradictory evidence does not become permission for broader inference, and thresholds are intended to be bounded and reviewable. The paper intentionally withholds reconstructive parameterization and anti-abuse internals.

7.4 Step 4: NICOLE evidence custody and release discipline

NICOLE Protocol governs evidence integrity, separation, evidence-custody metadata/state, and release. The objective is to preserve what happened during a qualified event without normalizing continuous remote custody of ordinary behavior. Incident evidence can therefore be more durable than ordinary-state telemetry while remaining more narrowly scoped.

7.5 Step 5: Degraded transport and fail-closed behavior

Transport logic attempts permitted heterogeneous paths rather than assuming a single uplink. Failure is not hidden. If communication cannot be completed, local state and evidence boundaries remain meaningful. This is central to the privacy thesis: resilience should not require constant external supervision. Fail-closed in this architecture primarily constrains unsupported authority, inference, and evidentiary claims; it does not require cessation of every otherwise-permitted protective communication.

7.6 Step 6: Human and institutional integration

LAKANA is not designed to replace emergency responders, incident command, medical professionals, or lawful institutional authority. Its purpose is to give those actors a bounded, interpretable safety packet when justified, while preventing the integration surface from becoming an ordinary-life monitoring portal.

8. Blue Force Bridge: Consent-Gated Rescue Communication

Blue Force Bridge (BFB) is the responder-facing communication concept inside the LAKANA Sovereign Operating System. Its design question is deliberately narrow: how can a responder receive enough trustworthy context to assist without acquiring permanent visibility into the civilian's ordinary life?

BFB treats responder access as a temporary safety event rather than a standing institutional entitlement. Requests are intended to be specific, separately authorized by data type, bounded by purpose and time, authenticated to an appropriate recipient, logged, and terminated when the protective purpose ends.

The public-safe rules below describe the intended authority boundary without exposing protected implementation logic. They also preserve an important distinction: forming, retaining, or transporting a safety packet is not the same as dispatch, responder action, or successful rescue.

Public-safe BFB access rules

Rule	Public-safe meaning
Specific request	A responder requests only the information category needed for the defined safety purpose.
User authority	When the user is able to respond, the user can accept or deny the request; denial is an auditable state, not a system failure.

Rule	Public-safe meaning
Unavailable user	User unavailability does not automatically create unlimited access. Any pre-authorized emergency rule must be explicit, bounded, and reviewable.
Separate data authorities	Audio, location, device status, and other evidence categories should not be treated as one undifferentiated permission.
Expiration	Access should terminate when the protective purpose terminates; standing ordinary-life visibility is not the default.
Evidence custody	Incident evidence may be preserved for integrity and later review without being automatically disclosed to every potential recipient.
Duress / coercion	A visible consent action may be unsafe under coercion; silent preservation, unsafe cancellation, and duress pathways require controlled validation.
Operational boundary	Packet formed, transport attempted, or evidence preserved does not mean dispatch occurred, an agency acted, or rescue was guaranteed.

BFB should therefore be evaluated as a controlled bridge into response systems, not as an open surveillance window. Human-factors testing must specifically evaluate consent comprehension, user unavailability, responder information sufficiency, access expiration, recipient authentication, and coercion/duress behavior.

9. Evidence Base: What the v71 DOI Supports

The companion LAKANA SOS v71 technical preprint, archived at DOI 10.5281/zenodo.19956214, reports the Final Core Connected Empirical Monte Carlo Run. The value of the evidence is not only the headline result. It is the combination of explicit run identity, trial counts, manifests, checksums, diagnostics, bounded comparators, and claim restrictions that allows an external reviewer to reconstruct the inferential chain without requiring public disclosure of protected implementation details [39].

The run used 250,000 full paired headline trials and 500,000 adversarial-stress trials. Under the modeled v71 assumptions, model-internal safe-state survivability was 0.975304 for LAKANA + CivOS, 0.973176 for LAKANA without CivOS, and 0.963996 for the modeled industry-style centralised/cloud comparator. The LAKANA + CivOS minus comparator difference was 0.011308, with a reported BCa bootstrap interval of [0.010908, 0.011732]. These are simulation quantities, not human survival probabilities.

Among the modeled architectural effects reported in v71, the largest reported separation was in degraded-communications transport resilience. Aggregate transport success was 0.939064 for LAKANA's heterogeneous transport model versus 0.674332 for the single-uplink comparator under the specified assumptions. TSARO and NICOLE diagnostics were also reported, including threat-conditioned escalation behavior and model-internal evidence-integrity measures. The technical manuscript is the controlling source for exact definitions, assumptions, exclusions, and statistical methods.

These results support one near-term conclusion: the architecture is sufficiently specified to justify adversarial review, hardware-in-the-loop testing, controlled degraded-communications experiments, and external replication planning. They do **not** establish field safety, certification, universal superiority, or procurement readiness.

v71 measure	Reported value	Permitted interpretation
Headline paired trials	250,000	Architecture-level Monte Carlo evidence under declared model assumptions.
Adversarial-stress trials	500,000	Stress testing of modeled failure/adversarial conditions.
Safe-state survivability: LAKANA + CivOS	0.975304	Model-internal outcome, not a real-world survival probability.

v71 measure	Reported value	Permitted interpretation
Safe-state survivability: cloud comparator	0.963996	Comparator within the specified reference architecture family.
Difference	+0.011308	Modeled difference; BCa interval [0.010908, 0.011732].
Transport success: heterogeneous LAKANA	0.939064	Modeled transport success under v71 channel assumptions.
Transport success: single-uplink comparator	0.674332	Reference transport result; motivates field validation.

10. Evidence-Tier Discipline and Claim Boundaries

LAKANA's public credibility depends on keeping evidence tiers explicit. A simulation can establish reproducibility, internal consistency, sensitivity, failure modes, and comparative behavior under declared assumptions. It cannot establish actual emergency outcomes, legal compliance, human-factors performance, RF performance in every environment, or independent cryptographic certification.

The claim level should therefore degrade when evidence inputs degrade. Full connected empirical-calibrated runs may support architecture-level simulation claims conditioned on those lanes. Synthetic/proxy fallback runs support engineering exploration and sensitivity analysis. Minimal deterministic runs support code-path verification. None should be relabeled as field evidence.

The modeled industry-style centralised/cloud comparator is also bounded. It represents a declared centralized/cloud reference architecture family, not every commercial product and not every possible privacy-enhanced cloud-edge hybrid. Broader comparator families are required before broader market claims are justified.

11. Falsifiable Research Questions and Validation Metrics

A protection-without-surveillance thesis is scientifically useful only if it can fail. LAKANA's central proposition would **not** be supported if controlled testing showed that reductions in continuous remote observation caused unacceptable degradation in threat handling, evidence completeness, emergency routing, or human usability relative to appropriately matched alternatives.

Conversely, evidence would support the proposition if LAKANA materially reduced non-event data egress, persistence, historical queryability, external custody, or inferential exposure while maintaining pre-registered thresholds for protective performance. The purpose of future pilots is therefore not to prove the company right; it is to test the architecture against declared failure criteria.

Research question	Example metric	Adverse result that would challenge LAKANA
Can non-event surveillance burden be reduced?	Bytes or state records leaving device during ordinary life; retention; queryability; number of custodians	No material reduction relative to comparator.
Can threat handling remain acceptable?	Threat-conditioned escalation; missed-event rate; false escalation	Materially worse safety discrimination after privacy reduction.
Can evidence remain useful?	Evidence-chain completeness; tamper detection; reviewer comprehension	Privacy architecture produces incomplete or unusable incident records.
Can communication survive degradation?	Delivery success; latency; route diversity; store-and-forward recovery	Local-first design fails more often than matched alternatives under realistic conditions.
Can users understand and control disclosure?	Consent comprehension; override rate; duress handling; time-to-action	Privacy controls create dangerous cognitive or response burden.

Research question	Example metric	Adverse result that would challenge LAKANA
Can responders use bounded data?	Responder task completion; information sufficiency; unnecessary-data rate	Responders require broad continuous visibility to achieve required outcomes.

12. Initial Pilot Domains

12.1 Rideshare, delivery, and mobile-worker safety

Drivers routinely work alone, at night, across unfamiliar locations, and in vehicles with strangers. A controlled non-emergency pilot can test user-triggered safety state, coercion-aware workflows, bounded evidence, transport degradation, and responder-facing packets without collecting every route as a permanent employer- or platform-owned location history.

12.2 HVAC, construction, and field-service crews

Field technicians enter private homes, roofs, crawlspaces, outage areas, and low-connectivity environments. Pilot endpoints can focus on check-in failure, device health, fall/immobility context, communication resilience, and bounded evidence while explicitly prohibiting generalized productivity surveillance.

12.3 Emergency-management and public-safety training contexts

Fire academies, emergency-management programs, and training cohorts can evaluate degraded-communications workflows, evidence usability, and incident reconstruction without claiming live dispatch or operational deployment. These environments are especially useful for testing whether responder needs can be met through purpose-bounded packets rather than standing access.

12.4 Campus, school, and youth-adjacent protection

Any pilot involving students or minors requires substantially higher governance thresholds. The initial research path should avoid broad student tracking and instead use adult-supervised, narrowly scoped training or research protocols with appropriate institutional review, parental permission, assent where applicable, minimization, retention limits, and explicit prohibition on disciplinary profiling.

13. Governance, Anti-Aggregation, and Non-Weaponization Requirements

LAKANA's ethical position is architectural rather than rhetorical. A system cannot merely promise not to misuse data; it must make unnecessary misuse harder by reducing what exists, who can access it, how long authority lasts, and what secondary purposes are technically permitted.

The central governance doctrine is **anti-aggregation**: ordinary safety operations should not create an ambient population-level mobility or behavior graph merely because storage and search make one technically possible. Cross-recipient access must be tied to purpose, authority, minimum necessary disclosure, expiration, and auditability. Aggregation that is necessary for a specific event or validated research protocol must be isolated from generalized person-level tracking and governed separately.

Non-weaponization follows from the same principle. LAKANA SOS is not designed as a law-enforcement dragnet, generalized employee-monitoring platform, advertising data source, insurance-scoring feed, or hidden behavioral eligibility system. This boundary does not mean public-safety institutions cannot receive useful information. It means the information should be event-bounded and justified rather than continuously available because the architecture collected everything by default.

- Data minimization by architecture: raw streams should not persist or leave the device merely because storage is available.
- Purpose limitation: evidence and access must be tied to a defined safety, legal, or research purpose.

- Consent and emergency-authority separation: ordinary consent, pre-authorized emergency release, unconsciousness, and duress require distinct rules.
- Expiration: responder or partner access should terminate when the protective purpose terminates.
- Anti-aggregation: ordinary non-event state must not silently become a generalized mobility or behavioral graph.
- No behavioral resale: safety telemetry must not be sold or repurposed into advertising, insurance pricing, employee scoring, or unrelated profiling.
- No hidden authority expansion: uncertain or contradictory evidence should not widen surveillance or decision authority.
- Evidence integrity: incident records should be hashable, separated, reviewable, and protected from silent alteration.
- External review: strong privacy, security, cryptographic, accessibility, or certification claims require independent evidence.
- Change governance: material changes to retention, sharing, or surveillance capability should require explicit review rather than silent feature drift.

14. Legal and Institutional Considerations

This paper does not offer a universal legal conclusion about ALPRs, location tracking, or safety telemetry. Constitutional doctrine differs between government and private actors, state constitutions may provide different protections, statutory regimes differ by context, and technological facts such as coverage, retention, precision, purpose, and access can alter legal analysis.

Carpenter, Jones, *Leaders of a Beautiful Struggle*, McCarthy, *NAACP v. Alabama*, and the ongoing Schmidt litigation should therefore be read as separate legal authorities illuminating different aspects of digital aggregation, location tracking, associational privacy, and persistent surveillance [19-24]. They do not establish one categorical rule that all persistent observation is unlawful.

For LAKANA pilots, institutional requirements will depend on context. Human-subject research may require IRB review and informed consent. School-associated data may trigger FERPA obligations. Health information may implicate HIPAA depending on the entities and relationships involved. Public-sector partners may have public-records, procurement, evidence-retention, CJIS, accessibility, cybersecurity, and incident-command requirements. Consumer deployments will require privacy, deceptive-practice, security, and state-law review. The system should be designed so these obligations narrow data and authority rather than become post hoc compliance layers placed over unrestricted collection.

Future LAKANA evaluation should also be anchored to established external engineering criteria rather than LAKANA-specific terminology alone. NIST SP 800-160 treats trustworthiness as a systems-engineering property across the lifecycle; the NIST AI Risk Management Framework emphasizes validity, reliability, safety, privacy, transparency, accountability, and harmful-bias management; and the NIST Cybersecurity Framework 2.0 provides governance and lifecycle cybersecurity structure [32-34]. These standards do not validate LAKANA. They provide independent criteria against which future controls, evidence, and operational claims can be tested.

Emergency-management integration likewise should fit existing coordination doctrine. FEMA's National Incident Management System and Community Lifelines framework emphasize interoperable coordination and continuity of essential services, while CISA's National Emergency Communications Plan and SAFECOM resources address resilient and interoperable emergency communications [35-38]. LAKANA SOS should be evaluated as a bounded civilian-side complement to those systems - not as a replacement for 911, public-safety radio, incident command, or emergency-management authority.

The legal minimum should not define the engineering optimum. A practice can be legally permitted and still impose unnecessary surveillance burden. LAKANA's objective is therefore to exceed bare compliance where doing so materially reduces risk without compromising the validated protective function.

15. Limitations

- The evidence remains architecture-level simulation evidence and public-safe technical disclosure. It does not establish field performance, deployment readiness, emergency-service certification, medical-device performance, law-enforcement suitability, or procurement readiness.

- The modeled industry comparator is a centralized/cloud reference architecture family, not every possible commercial system or privacy-enhanced hybrid.
- v71 includes empirical-calibrated and proxy-conditioned lanes; proxy inputs do not substitute for hardware-captured field evidence.
- Transport results are promising within the model but real RF performance varies with geography, construction materials, weather, congestion, interference, device power, and regulatory constraints.
- TSARO escalation is meaningful but imperfect; false escalation, missed threat conditions, alert fatigue, human workload, and adversarial manipulation require controlled testing.
- NICOLE simulation diagnostics do not equal independent cryptographic certification, evidentiary admissibility, or privacy/security audit.
- Blue Force Bridge requires substantial consent, duress, recipient-authentication, accessibility, and human-factors validation before live use.
- The proposed surveillance-burden framework is a research scaffold, not a validated universal privacy metric. Weights and thresholds must be pre-registered for each study.
- Legal analysis in this paper is explanatory and not legal advice. Court decisions are context-specific, and several contemporary ALPR questions remain unresolved.
- The architecture intentionally withholds reconstructive implementation details, protected thresholds, and anti-abuse internals. This limits full public reproduction and requires controlled-access review for some claims.

16. Validation and Translation Roadmap

The next stage should convert internally generated architecture and simulation evidence into independently reviewed and independently reproduced evidence. Publication is one component of that transition, but publication alone is not validation. The appropriate sequence is external criticism, bench measurement, controlled pilot design, independent analysis, and progressively stronger field claims only when the evidence supports them.

The commercialization model should remain subordinate to that evidence ladder. Non-dilutive funding, university partnerships, emergency-management relationships, and cloud infrastructure should be used to finance validation, audit, reproducibility, hardware prototypes, and secure evidence infrastructure - not to force broad deployment before the architecture has earned stronger claims.

Stage	Primary work	Evidence earned
1. Public record	SSRN working paper; Zenodo technical preprint; provenance and claim boundaries	Discoverability and stable research record, not validation.
2. External review	Privacy/security review; statistical review; responder/human-factors critique	Independent identification of defects and testable claims.
3. Bench / HIL	Transport, power, latency, failure injection, evidence integrity	Measured device/system performance under controlled conditions.
4. Controlled pilot	Adult volunteer or training cohorts; explicit consent; narrow scenarios	Human usability and field-adjacent evidence within defined scope.
5. Independent replication	University/lab or external evaluator repeats defined protocols	External credibility and calibration updates.
6. Deployment-readiness	Security/privacy audit; governance review; partner-specific controls	Basis for bounded operational claims; still not universal certification.

17. Conclusion: Protection Without Persistent Observability

The future of civilian safety should not be reduced to a choice between being protected and being observable. Searchable historical data can create retrospective reconstruction power; persistent live visibility can create interception power; and behavioral aggregation can create inferential power. Each capability may serve legitimate purposes, but each also creates attack surface, coercive potential, function-creep risk, and governance obligations that must be evaluated alongside protective performance.

LAKANA SOS - the LAKANA Sovereign Operating System for civilian safety - is one proposed architecture for testing whether those powers can be reduced without unacceptable loss of useful safety capability. It keeps safety state local where feasible, reduces raw observations toward bounded evidence, fails closed under critical uncertainty, preserves event evidence through NICOLE, constrains institutional access through event-scoped authority, and treats degraded communication as a resilience problem rather than a reason for permanent cloud supervision.

The proposition remains falsifiable. If materially lower surveillance burden produces unacceptable losses in threat handling, evidence integrity, transport resilience, responder usability, or human safety, the architecture must be revised or rejected. If comparable declared protective-performance requirements can be satisfied while materially reducing ordinary-life data egress, persistence, queryability, external custody, and inferential exposure, the lower-burden architecture earns stronger support.

Many data-intensive safety architectures focus primarily on controlling access after broad observation has occurred. LAKANA SOS asks an earlier engineering question: what ordinary-life data must exist at all for the declared protective function?

Statements and Declarations

Data and Code Availability. This white paper is a public-facing companion to the LAKANA SOS v71 technical preprint archived at DOI 10.5281/zenodo.19956214. The public record includes aggregated results, figure-level summaries, evidence-tier descriptions, claim boundaries, and a public-safe interpretation of the architecture. Full implementation details, protected thresholds, anti-abuse logic, private chain-of-custody details, and reconstructive engine internals are withheld for intellectual-property, safety, and misuse-prevention reasons. Controlled-access review may be pursued under appropriate confidentiality, legal, and safety terms.

Competing Interests. MarTaize K. Fails is the founder of LAKANA Sovereign Systems LLC and the inventor/developer of the evaluated architecture.

Funding. No external funding supported this white paper.

Ethics. This white paper reports architecture, simulation evidence, legal/privacy analysis, and validation planning. It does not report human-subject deployment, clinical intervention, emergency-response deployment, or field trial results.

AI Assistance. OpenAI ChatGPT was used for author-directed drafting, source synthesis, editing, formatting, and publication preparation. The author selected the sources and claims for inclusion, directed revisions, reviewed the final manuscript, and remains responsible for the content.

References

- [1] Jon Keegan, “The Popular Family Safety App Life360 Is Selling Precise Location Data on Its Tens of Millions of Users,” The Markup (Dec. 6, 2021). [The Markup article](#)
- [2] Jon Keegan, “Life360 Says It Will Stop Selling Precise Location Data,” The Markup (Jan. 27, 2022). [The Markup article](#)
- [3] Office of the Texas Attorney General, “Attorney General Ken Paxton Sues Allstate and Arity...” (Jan. 13, 2025). [Texas OAG release](#)
- [4] State of Texas v. Allstate Corp. et al., Original Petition (Montgomery County, Texas, Jan. 13, 2025). [Texas petition \(PDF\)](#)
- [5] Federal Trade Commission, X-Mode Social, Inc. / Outlogic matter and final order (2024). [FTC case page](#)
- [6] Federal Trade Commission, “FTC Takes Action Against Gravy Analytics, Venntel...” (Dec. 3, 2024); final order Jan. 14, 2025. [FTC release](#)
- [7] Federal Trade Commission, “FTC Takes Action Against Mobilewalla for Collecting and Selling Sensitive Location Data” (Dec. 3, 2024). [FTC release](#)
- [8] Federal Trade Commission, “FTC Says Ring Employees Illegally Surveilled Customers, Failed to Stop Hackers...” (May 31, 2023). [FTC release](#)
- [9] U.S. Department of Justice, Eastern District of Pennsylvania, Victor Ortiz stalking / GPS-tracker case; sentencing release (Apr. 11, 2025). [DOJ release](#)
- [10] U.S. Attorney’s Office for the District of Columbia, Ulus Jorden GPS-stalking case (Jan. 2012). [DOJ release](#)
- [11] U.S. Department of Justice, Western District of New York, “Former Irondequoit Police Officer Sentenced For Cyber Stalking His Ex-Girlfriend” (Feb. 1, 2018). [DOJ release](#)
- [12] U.S. Department of Justice, Eastern District of New York, “Two Men Charged for Breaching Federal Law Enforcement Database...” (Mar. 14, 2023). [DOJ release](#)
- [13] Federal Trade Commission, SpyFone matter (2021), including ban from surveillance business. [FTC release](#)
- [14] Federal Trade Commission, General Motors / OnStar final order regarding geolocation and driving behavior data (Jan. 14, 2026). [FTC release](#)
- [15] Federal Trade Commission, BetterHelp final order regarding sensitive health data and advertising (July 14, 2023). [FTC release](#)
- [16] Federal Trade Commission, Avast final order regarding browsing-data sales (June 26, 2024). [FTC release](#)
- [17] U.S. Department of Justice, United States v. Meta Platforms, Inc., Fair Housing Act algorithmic-advertising settlement (2022). [DOJ case page](#)

- [18] Federal Trade Commission, Rite Aid facial-recognition matter (Dec. 19, 2023). [FTC release](#)
- [19] Carpenter v. United States, 585 U.S. 296 (2018). [Supreme Court opinion \(PDF\)](#)
- [20] United States v. Jones, 565 U.S. 400 (2012). [Justia case page](#)
- [21] Leaders of a Beautiful Struggle v. Baltimore Police Department, 2 F.4th 330 (4th Cir. 2021) (en banc). [Fourth Circuit opinion \(PDF\)](#)
- [22] Commonwealth v. McCarthy, 484 Mass. 493 (2020). [Massachusetts SJC / Justia](#)
- [23] NAACP v. Alabama ex rel. Patterson, 357 U.S. 449 (1958). [Supreme Court / Justia](#)
- [24] American Civil Liberties Union, Schmidt v. Norfolk case page (ongoing; updated Apr. 20, 2026). [ACLU case page](#)
- [25] Flock Safety, “Law Enforcement Data Access,” Trust Center (accessed Sept. 7, 2026). [Flock Trust Center](#)
- [26] Flock Safety, “Flock Implements Enhanced Guardrails Across California...” (Mar. 2, 2026; updated July 24, 2026). [Flock guardrails disclosure](#)
- [27] Flock Safety, Press Center / retention and misuse safeguards (accessed Sept. 7, 2026). [Flock Press Center](#)
- [28] Matthew Gault, “The Mayor Who Said ‘No’ When Flock Came to Town,” 404 Media (Aug. 10, 2026). [404 Media report](#)
- [29] KOCO 5, “City Council approves contract renewal to keep using Flock cameras in Oklahoma City” (Aug. 18, 2026). [KOCO report](#)
- [30] KOCO 5, “As OKC votes to keep Flock cameras, police chief reveals 109 agencies have access to data” (Aug. 18, 2026). [KOCO report](#)
- [31] Reuters, “Florida bans highway license-plate readers as backlash over surveillance spreads” (Sept. 3, 2026). [Reuters report](#)
- [32] Ross, R., Winstead, M., & McEvelley, M., Engineering Trustworthy Secure Systems, NIST SP 800-160 Vol. 1 Rev. 1 (2022). [NIST SP 800-160](#)
- [33] NIST, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (2023). [NIST AI RMF DOI](#)
- [34] NIST, Cybersecurity Framework 2.0, NIST CSWP 29 (2024). [NIST CSF 2.0 DOI](#)
- [35] Federal Emergency Management Agency, National Incident Management System, 3rd ed. (2017). [FEMA NIMS \(PDF\)](#)
- [36] Federal Emergency Management Agency, Community Lifelines Implementation Toolkit 2.1 (2023). [FEMA Lifelines Toolkit \(PDF\)](#)
- [37] Cybersecurity and Infrastructure Security Agency, National Emergency Communications Plan. [CISA NECP](#)
- [38] Cybersecurity and Infrastructure Security Agency, SAFECOM. [CISA SAFECOM](#)
- [39] Fails, MarTaize K., LAKANA SOS v71 technical preprint, Zenodo DOI 10.5281/zenodo.19956214 (2026). [Zenodo DOI](#)
- [40] Flock Safety, License Plate Reader Policy (accessed Sept. 7, 2026). [Flock LPR Policy](#)