# LAKANA

## Sovereign Safety System

*Privacy-Preserving Civilian Safety Infrastructure for Protection Without Surveillance*

**MarTaize K. Fails**
Founder, LAKANA Sovereign Systems
Norman, Oklahoma, USA • May 2026
www.lakana.systems
*Public White Paper / SSRN Working Paper v1.1*

CORE THESIS

Civilian protection should not require continuous surveillance, data brokerage, or cloud dependence. LAKANA SOS is designed around local-first safety state, deterministic fail-closed logic, evidence custody, consent-gated response, degraded-transport resilience, and staged institutional validation. Safety and sovereignty are not a trade-off. They are the same requirement.

**Companion technical preprint:** LAKANA SOS v71 Final Core Connected Empirical Monte Carlo Run. DOI: 10.5281/zenodo.19956214

SCOPE STATEMENT

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

**ABSTRACT**

LAKANA Sovereign Safety System is a proposed privacy-preserving civilian safety infrastructure centered on LAKANA SOS: a local-first, fail-closed cyber-physical safety operating system for degraded, adversarial, and infrastructure-fragile conditions. The design responds to a structural failure in many consumer and field-safety products: protection is coupled to continuous location visibility, cloud dependence, and downstream data brokerage. Public enforcement actions by the FTC, state attorneys general, and federal litigation document instances in which safety-branded applications became surveillance supply chains, exposing users — including minors — to data-broker sales, insurance-rate scoring, and downstream AI inference without meaningful consent.

LAKANA reverses that premise architecturally. Safety state is kept local by default. Escalation is governed by the TSARO risk-envelope engine, which applies deterministic, evidence-constrained state transitions rather than probabilistic expansion. Incident evidence is separated, hashed, and governed by the NICOLE protocol before any external release. Responder access is delivered through Blue Force Bridge, a consent-airlock architecture providing rescue-scoped access with ephemeral session isolation and a complete audit trail. Transport is treated as a survivability layer with multiple permitted fallback paths; failed attempts are recorded as explicit state.

The companion technical preprint — LAKANA SOS v71 Final Core Connected Empirical Monte Carlo Run, archived at DOI 10.5281/zenodo.19956214 — reports 250,000 full trials and 500,000 adversarial-stress trials under empirically calibrated parameters. Under the modeled assumptions, LAKANA SOS + CivOS achieved higher simulated safe-state survivability than a modeled industry-style centralized/cloud comparator, with transport success of 0.939 versus 0.674 for the comparator. These results constitute architecture-level simulation evidence, not field validation. This white paper presents the problem, surveillance-risk landscape, architecture, Blue Force Bridge, simulation evidence posture, institutional fit, and validation roadmap for LAKANA SOS.

**KEYWORDS**

*local-first safety • fail-closed systems • civilian protection • privacy-preserving safety infrastructure • safety without surveillance • location privacy • data brokerage • TSARO risk envelope • NICOLE protocol • Blue Force Bridge • degraded infrastructure • consent architecture • Monte Carlo simulation*

# Executive Summary

The civilian safety market has a structural contradiction at its center. Products marketed around care, family protection, crash detection, and driver safety often depend on continuous location extraction and downstream behavioral data flows. That data does not stay in one place. It moves to analytics firms, data brokers, insurance subsidiaries, advertising exchanges, and AI inference engines, creating a surveillance chain the user never consented to and cannot exit. The documented consequences range from federal litigation to state enforcement actions to insurance-rate changes affecting tens of millions of Americans.

LAKANA SOS inverts that architecture entirely. Safety begins with user sovereignty. The system permits bounded safety actions only when local evidence, state integrity, and governed release rules justify them. Raw telemetry stays on the device. Escalation is deterministic and bounded. Evidence is hash-sealed before any external delivery. Responder access is rescue-scoped and logged. A failed network path is recorded as state, not erased as absence.

The system has six architectural functions: (1) local safety state, allowing the device to make bounded decisions without a cloud service; (2) deterministic risk envelopes through TSARO, which restricts actions to states permitted by evidence rather than expanded by probabilistic confidence; (3) evidence custody through NICOLE, which records, separates, hashes, and governs incident records; (4) degraded transport, attempting multiple permitted delivery paths and recording success or failure across all channels; (5) coercion-aware release, distinguishing ordinary consent from duress, unconsciousness, and unsafe disclosure; and (6) an institutional audit path, making the architecture reviewable by external advisors, grant agencies, emergency management partners, and validators.

The companion technical preprint — LAKANA SOS v71 Final Core Connected Empirical Monte Carlo Run, archived at DOI 10.5281/zenodo.19956214 — reports 250,000 full trials and 500,000 adversarial-stress trials. Under the modeled assumptions, LAKANA SOS + CivOS achieved higher simulated safe-state survivability than a modeled industry-style centralized/cloud comparator, with transport success of 0.939 versus 0.674 for the comparator. Those findings support pilot validation and non-dilutive funding pursuit. They do not replace field validation.

The immediate commercialization path is advisor-first. LAKANA requires institutional advisors and partner validation before large pilots: SBIR/STTR advisors for agency fit and proposal discipline, university partners for study design and privacy review, emergency management advisors for operational fit, and cyber/privacy advisors for abuse-resistance assessment.

## 1. The AI Surveillance Default: You Pay With Everything

There was a time when surveillance required a warrant, a dedicated resource, and a named target. That era is over. The architecture of the connected internet has made continuous behavioral capture the ambient condition of digital life. You do not opt into it. You absorb it the moment you use any connected application. What has changed in the past three years is not the existence of this data extraction — it is what that data can now do.

AI has made the transition qualitatively worse. A decade ago, location data showed a dot on a map. Today that same data feeds inference engines that deduce income bracket, health status, religious practice, relationship stress, political alignment, employment risk, and domestic vulnerability — all without a single direct disclosure. The same telemetry that once showed a delivery route now reveals a home address, a school, a clinic visit, a custody exchange location, a late-night labor pattern, a domestic escape route. Precision location data is not a neutral signal type. It is a window into everything a person might reasonably wish to keep private.

The implicit bargain of the AI era has become explicit and unavoidable: give us everything about yourself, and we will give you a useful tool. The companies offering safety applications are not exempt from this dynamic. In many cases, they are among its most significant participants — because safety requires presence, and presence generates the richest possible behavioral signal.

**THE SURVEILLANCE TRAP**

Ambient behavioral data — location, motion, routine, association — has become the raw material that trains models, targets advertising, scores insurance risk, and predicts purchasing behavior.

> Users are not customers of these systems. They are the product. The only durable protection is an architecture that governs before it collects, because data that is never extracted cannot be sold, subpoenaed, breached, or weaponized.

This is not a theoretical concern. Federal agencies, state attorneys general, and plaintiff law firms have all documented specific instances in which safety-branded applications became surveillance supply chains. The following section details those cases with precision — because the only way to understand why LAKANA's architecture matters is to understand what the current architecture has already done to real people.

# 2. The Data-Broker Safety Model: Documented Cases

The following cases are drawn from public reporting, federal enforcement orders, and court filings. They are not presented to single out individual companies as uniquely malicious. They document a structural pattern: safety applications built on a surveillance business model inevitably expose users in ways that contradict their stated purpose. The pattern is not an accident. It is the output of an architecture that collects first and governs afterward, if at all.

## 2.1 Life360: When a Family Safety App Becomes a Data Supply Engine

Life360 is among the most widely recognized family-safety applications in the United States, marketed to parents as a way to know their children's whereabouts. In December 2021, The Markup published an investigation documenting that Life360 was selling the precise location data of its users — including minors — to multiple data brokers simultaneously. Named brokers included SafeGraph, Cuebiq, X-Mode, and Arity. The reporting described Life360 as one of the largest commercial sources of precise location data in the country, with the practice generating tens of millions of dollars per year in revenue.

Life360 disclosed the data-sharing in its privacy policy. A disclosure embedded in a document users are not expected to read, negotiated on terms users cannot meaningfully reject if they want the product's safety features, is not meaningful consent. It is the structural appearance of consent without the substance. Following the investigation, Life360 announced it would stop selling precise data to most brokers. The Markup subsequently reported that arrangements with Arity and aggregated data flows to Placer.ai would continue.

The lesson is architectural: a safety product becomes a data-supply engine the moment its business model requires it to. Privacy policy language does not prevent this. A different data model does.

## 2.2 E.S. v. Life360 Inc.: Federal Litigation Over Children's Location Data

A proposed class action, E.S. v. Life360 Inc. (No. 4:23-cv-00168, N.D. Cal., filed January 12, 2023), alleged that Life360 sold users' precise geolocation data — including data concerning minors — without adequate consent. The complaint named SafeGraph, Arity, Cuebiq, and X-Mode as downstream recipients. The case was later voluntarily dismissed with prejudice; the allegations were not adjudicated.

The case remains significant as a structural risk signal. Family-safety data involving minors is sensitive enough that its alleged commercial transfer to advertising and analytics firms produced federal litigation. Whether or not the specific allegations would have been proven, the data chain they describe — safety app, data broker, analytics firm, downstream purchaser — is exactly the architecture LAKANA is designed to prevent. Not through policy promises. Through a different data model.

## 2.3 Texas v. Allstate and Arity: Insurance Scoring from Covert Mobile Telemetry

In January 2025, the Texas Attorney General filed suit against Allstate and its subsidiary Arity, alleging unlawful collection, use, and sale of location and movement data from Texas residents' cell phones through allegedly secretly embedded software in third-party mobile applications. Among the apps named in related public reporting were Life360, Fuel Rewards, GasBuddy, and Routely. The State of Texas alleged that Allstate had amassed behavioral data on more than 45 million Americans through mobile-app developer relationships.

The core allegation is that driving behavior collected through these applications was used to justify insurance-rate adjustments, meaning a user's premium was being set, at least in part, by behavioral data extracted from an application installed for an entirely different purpose. Allstate has denied the allegations and asserted consent-based defenses. Reuters reported in March 2026 that related private litigation survived early dismissal.

The structural pattern the allegations identify is the point: a sensor-rich application, a data-licensing relationship with an insurance analytics subsidiary, behavioral scoring as an output, and insurance-rate consequences the user could not anticipate or contest. That pipeline is not unique to one company. It is a business model. LAKANA is an architectural argument against it.

## 2.4 Federal Trade Commission: Four Enforcement Actions in Two Years

The FTC has pursued a sustained enforcement effort against the commercial location-data industry that documents a sector-wide pattern rather than isolated wrongdoing.

The FTC's action against X-Mode Social (operating as Outlogic) centered on allegations that precise location data could track visits to medical and reproductive health clinics, places of worship, and domestic-abuse shelters. The consent order, finalized in April 2024, prohibited the sale of sensitive location data and required deletion of certain previously collected data sets.

The FTC's litigation against Kochava alleged the sale of geolocation data from hundreds of millions of devices, with the complaint specifically citing the ability to trace movements to reproductive health clinics, shelters, and addiction-recovery facilities. In January 2025, the FTC finalized orders against Gravy Analytics and Venntel, prohibiting them from selling or sharing sensitive location data. In December 2024, the Commission acted against Mobilewalla for alleged collection and sale of sensitive location data enabling demographic inference.

The pattern across four actions in two years is consistent: precision location data, collected through SDK embedding or app developer relationships, enriched with demographic inference, and sold to

downstream purchasers making consequential decisions about real people. These actions are not anomalies. They are a cross-section of a market structure.

**THE STRUCTURAL LESSON**

These enforcement actions share a common architecture: a safety or convenience application collects data, embeds a third-party SDK or sells developer access, and the resulting data flows to brokers, insurers, advertisers, and AI-inference platforms.

The user's consent to the original application does not constitute consent to the downstream chain. No policy revision prevents this. The only durable protection is an architecture that governs before it collects — because data that is never accumulated cannot be sold, subpoenaed, breached, or weaponized.

## 2.5 The Architecture That Changes the Outcome

The failure documented above is a design failure, not a policy failure. The legacy location-safety model begins with continuous visibility and attempts to govern it afterward. LAKANA begins with user sovereignty and permits safety actions only when bounded evidence justifies them. The figure below summarizes the structural difference.

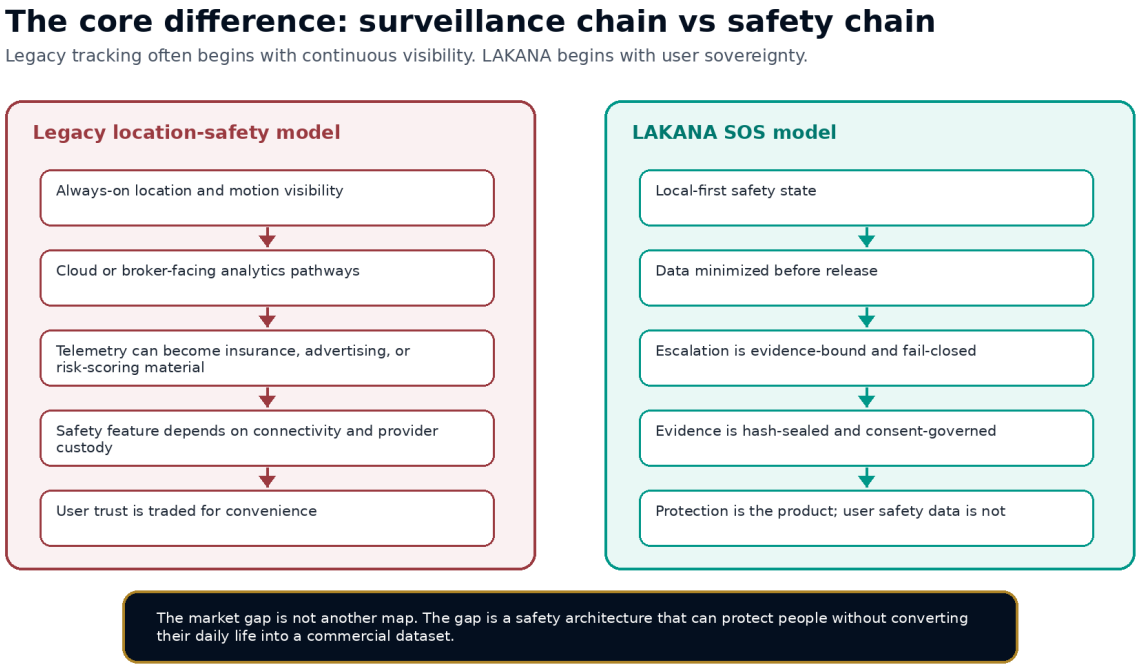


*Figure 1. The structural difference: continuous-visibility model versus evidence-bound local-protection model.*

**OPERATING PRINCIPLE**

No safety action is permitted merely because a model is confident. The action must be justified by bounded evidence, local state, user-governed release rules, and the audit chain. A safety

architecture that cannot be independently audited cannot be trusted.

# 3. What LAKANA SOS Is

LAKANA SOS is civilian safety infrastructure designed to operate as close to the user and the physical event as possible. It is built for protection under ordinary conditions and under degraded conditions where a conventional application may lose cloud access, cellular transport, GPS confidence, or evidence custody. It does not assume the internet is always present. It does not assume a remote service can always decide. It does not assume all sensor readings are honest. It does not assume a user under coercion can safely interact with a screen in the normal way.

The system is not a replacement for 911, emergency managers, incident command, medical devices, law enforcement, or professional responders. It is a civilian-side safety substrate that preserves evidence, enables bounded escalation, and supports partner workflows when formal response systems are not yet engaged or when infrastructure is degraded.

## The Six-Step Architectural Sequence

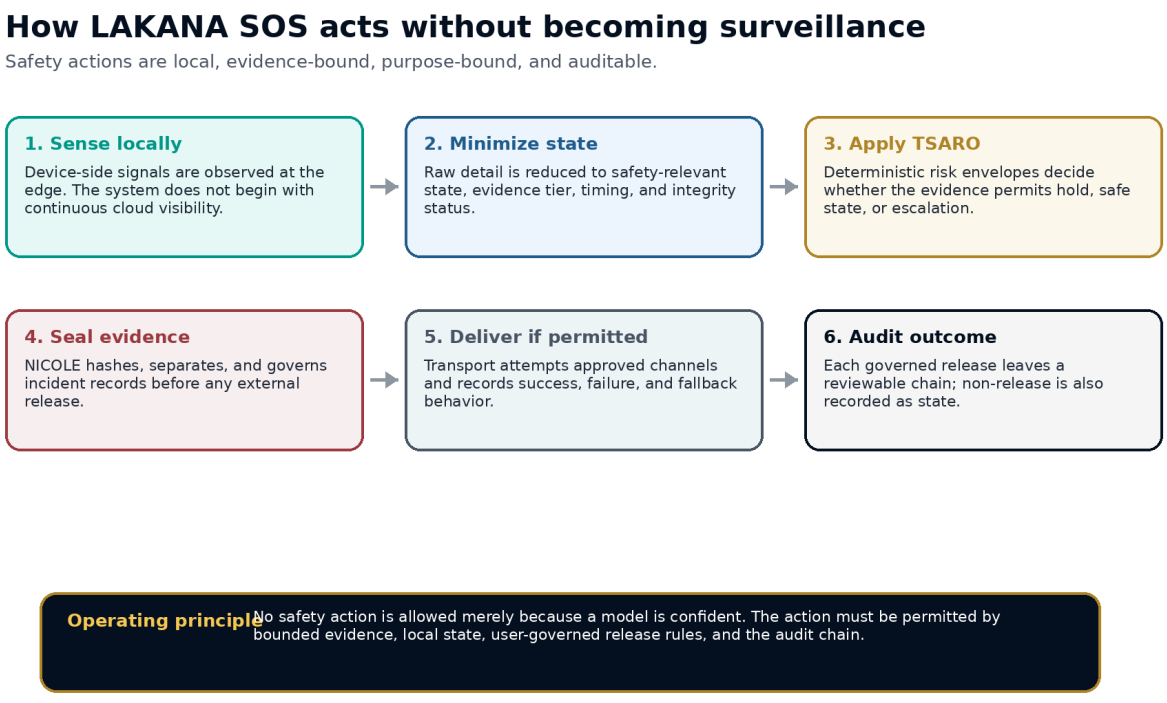


*Figure 2. Architectural sequence: local observation, state minimization, TSARO envelope, NICOLE evidence sealing, governed transport, audit closure.*

| Step | Action | What the User or Institution Gains |
|---|---|---|

| | | |
|---|---|---|
| 1 | **Sense Locally**<br>Device-side signals are evaluated on or near the user's device. The system does not begin with continuous cloud streaming. | Protection functions even when a cloud path is unavailable. Raw data exposure is minimized from the first moment. |
| 2 | **Minimize State**<br>Signals are converted into bounded state, evidence tier, timestamp, and integrity status. Raw detail is retained only when an incident pathway justifies it. | The user is not required to trade constant personal visibility for safety. |
| 3 | **Apply TSARO**<br>The risk layer checks whether evidence permits no action, hold, local safe state, escalation, or preservation. Contradiction narrows behavior — it never expands it. | Actions are tied to permitted states, not opaque probabilistic guesses. The system is auditable. |
| 4 | **Seal Evidence**<br>NICOLE hashes, separates, and governs incident records before any external release. Release is purpose-bound and logged. | Post-incident records are harder to alter, erase, or weaponize without leaving a governance trace. |
| 5 | **Deliver if Permitted**<br>Transport attempts multiple approved channels. Success means a safety-relevant packet found at least one permitted route. | A single failed network path does not collapse alerting, logging, and responder coordination. |
| 6 | **Audit Outcome**<br>Each governed release leaves a reviewable chain. Non-release is also recorded as state. | Advisors, partners, and external reviewers can inspect outcomes without open access to the full implementation. |

# 4. Architectural Components

The subsystem descriptions below disclose architecture, purpose, and design logic at the level required for institutional evaluation. Threshold ladders, internal coefficients, key schedules, and trade-secret implementation details are withheld. All configuration parameters needed to interpret simulation results are disclosed in the companion technical preprint at DOI 10.5281/zenodo.19956214.

## 4.1 TSARO: Deterministic Risk-Envelope and Skepticism Layer

TSARO is the safety decision layer. Its job is not to guess harder. Its job is to determine whether the current evidence permits a safety action at all. A probabilistic model may say an event is likely. TSARO asks a stricter question: does the local state satisfy the permitted envelope for action, release, hold, or suppression?

TSARO performs four steps. It receives minimized local state and evidence-tier labels. It checks consistency across available signals and context. It applies bounded transition rules. It emits a permitted system state: no action, hold, local safety preparation, evidence preservation, or governed escalation. When evidence is contradictory, missing, stale, or below the envelope, TSARO does not expand authority. It holds or fails closed. Contradiction is not ambiguity to be resolved by confidence. It is a veto.

For a lay reader, this means LAKANA does not panic every time a phone drops or a GPS point jumps. For a technical reviewer, it means the safety loop is constrained by deterministic state-machine behavior rather than unconstrained stochastic inference. For a grant reviewer, it means the architecture has a clear validation target: measure whether TSARO improves true threat response without inflating false escalation and alert fatigue.

## 4.2 NICOLE: Evidence Custody and Governed Release

NICOLE is the evidence layer. It does not decide whether a person is safe. It protects the integrity of the records that explain what happened. In a conventional safety application, logs may be server-side, mutable, incomplete, provider-dependent, or simply unavailable when the network fails. LAKANA treats evidence as a safety object: capture is minimized, separation is enforced, release is governed, and later review must be possible regardless of network conditions at the time of the incident.

The NICOLE sequence is direct. A qualifying event creates a minimized record. The record is timestamped and hashed. Related streams are separated so one channel cannot silently contaminate another. Release status is checked against consent, safety purpose, and configured authority. Permitted release creates an auditable trail. Non-release also creates a recordable state — because a withheld action can be as important as an action taken.

Evidence custody is a trust product. A school, city, rideshare organization, HVAC crew, emergency management partner, or insurer does not only need an alert. It needs to know whether the alert was justified, what was released, what was withheld, whether consent boundaries were respected, and whether the record can be audited after the fact. NICOLE provides that chain without requiring continuous behavioral surveillance as its substrate.

## 4.3 CivOS: Local Survival Substrate

CivOS is the local survival substrate that supports power-state awareness, persistent state, transport health, and degraded execution. Safety logic cannot be trusted if the substrate beneath it is silently failing. Battery state, operating-system health, persistent record writing, and transport availability are safety variables, not background conveniences.

CivOS makes degraded substrate conditions visible to the safety path. If the device is low on power, if persistent state cannot be written, if transport routes are failing, or if the local execution environment becomes unreliable, the system does not behave as though everything is normal. Degradation

becomes explicit state. That is the difference between a safety system that fails silently and one that fails closed.

### 4.4 Transport: What Transport Success Actually Means

Transport success does not mean the internet is working. It means that a safety-relevant event packet, beacon, evidence-status notice, or configured responder message found at least one permitted delivery path under the modeled conditions. In a legacy single-uplink application, one failed path can collapse alerting, logging, and escalation together. In LAKANA SOS, transport is treated as a survivability layer with multiple permitted channels, logged outcomes, and no silent erasure of failed attempts.

This matters for rural routes, rideshare drivers, field technicians, storm-damaged neighborhoods, building interiors, and event venues — precisely the environments where cellular fails at the exact moment a user needs help. The companion technical preprint at DOI 10.5281/zenodo.19956214 reports empirically calibrated transport success of 0.939 for the LAKANA architecture versus 0.674 for the modeled industry-style comparator. That result is a pilot target, not a field guarantee. It is a strong reason to make degraded-communications validation the first serious field test.

## 5. Blue Force Bridge: Consent-Gated Responder Access

Every safety system eventually faces the same hard problem: the moment a person needs help most is also the moment they may be least able to grant it. They may be unconscious, incapacitated, coerced, or in a physical environment where any visible action on their device could escalate the threat. The standard industry response to this problem has been to resolve the tension by defaulting to access — giving responders, platforms, or monitoring services broad visibility into user data on the assumption that availability protects life. LAKANA rejects that resolution as architecturally unsound. Broad default access does not protect the user. It transfers risk from the emergency scenario to every scenario that follows — including the ones where the same access architecture is exploited by an abuser, a coercive institution, a data broker, or a hostile party who controls the device. Blue Force Bridge is the answer to that hard problem. It is the architectural boundary that allows a legitimate, verified responder to receive exactly what a safety event requires, under governed terms, with a complete audit trail, without that access becoming a permanent open channel into the user's life.

The core design principle is rescue-scoped access: a responder receives what is necessary for the active safety event, and nothing more. This is not a policy promise. It is an architectural constraint enforced at the protocol level. A request that does not specify a safety function type, a declared scope, and a governing authority cannot be fulfilled. A session that was granted for scene stabilization cannot be expanded to retrieve movement history, prior audio, or behavioral data without a new request cycle that passes through the full consent-airlock sequence. The reason this matters is not just legal hygiene. It is operationally sound. A legitimate responder in a structural collapse, a rideshare incident, or a field-work accident does not need general surveillance access to do their job. They need breathing status. They need last-known location. They need scene context. They need to be able to speak to the person they are trying to reach. Bounded access gives them all of that. Unbounded access gives them liability, evidentiary contamination risk, and a data surface that can be subpoenaed, breached, or misused long after the incident is resolved.

**DESIGN RULE**

Rescue-scoped access is permitted only when it serves the user's specific safety event and leaves a complete audit trail. General-purpose listening, background monitoring, or responder access to user history is outside the architectural boundary. Bounded access with an audit trail is operationally more useful to a legitimate responder than open access with no chain of custody — because the responder operates in a professional context where documentation and post-incident accountability are not optional.

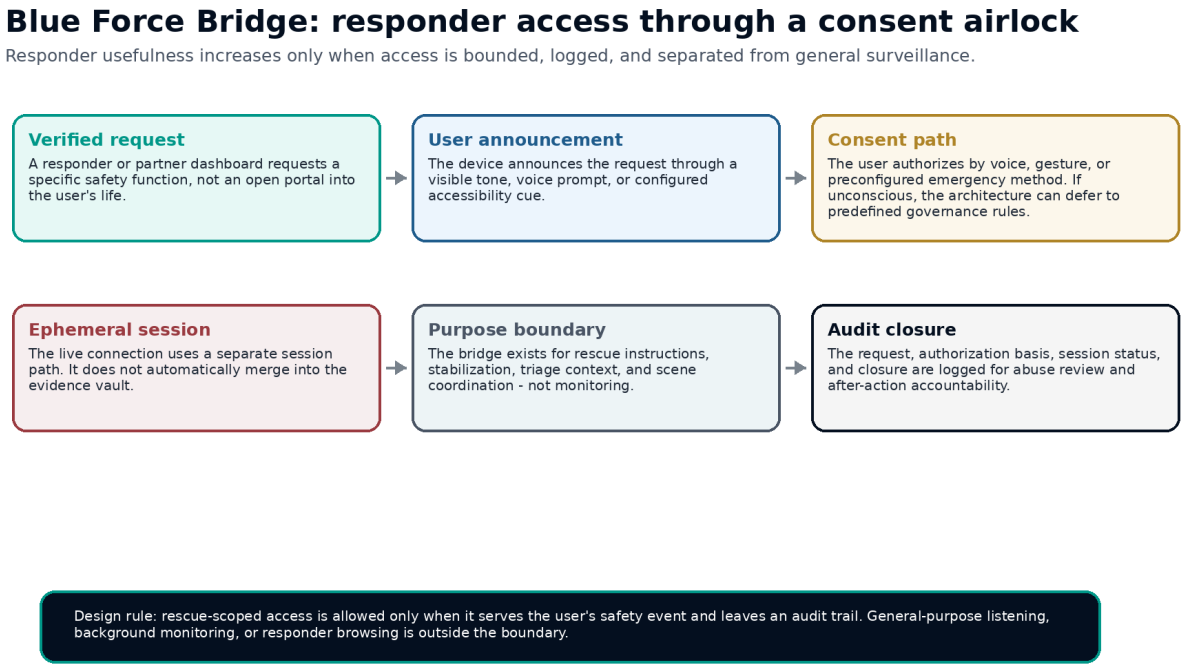


*Figure 3. Blue Force Bridge: every responder access request is typed, scoped, announced, authorized, and logged. No access is anonymous, unaudited, or scope-unlimited.*

## The Consent Airlock Sequence

The consent airlock is not a user-interface design choice. It is a trust architecture. Each stage in the sequence serves a distinct safety and accountability function, and no stage can be skipped by a higher-privilege caller. The architecture does not have an administrative override that bypasses the airlock. If the user is unable to respond and the event qualifies for a predefined emergency governance override, that override is itself a declared, logged state — not a silent bypass. Every access event, whether granted, deferred to an override, or denied, becomes a permanent reviewable entry in the audit chain.

| # | Stage | What Happens and Why It Matters |
|---|---|---|
| 1 | **Verified Request** | A responder or partner dashboard requests a specific safety function — not an open portal. The request is typed, scoped, and logged before any response begins. |

| | | |
|---|---|---|
| 2 | **User Announcement** | The device announces the request through a configured voice prompt, tone, screen cue, or accessibility signal. The user is informed before any access begins. |
| 3 | **Consent Path** | The user authorizes through a voice keyword, emergency gesture, or preconfigured method. If the user is unconscious or unable to respond, the system defers to predefined emergency governance rules, not a default open grant. |
| 4 | **Ephemeral Session** | The live connection uses a separate ephemeral channel and separate key material. It does not automatically merge into the evidence vault. A responder intercom and a forensic evidence pack are distinct objects with distinct risk profiles. |
| 5 | **Purpose Boundary** | The bridge exists for rescue stabilization, triage context, and scene coordination only. The session scope is defined at initiation and cannot expand without a new request cycle. |
| 6 | **Audit Closure** | The request, authorization basis, session status, and closure are logged for post-incident accountability. Both granted and denied access become reviewable state. |

The sequence is designed to prevent two distinct failure modes simultaneously. The first is unauthorized access — a hostile party, a coercive institution, or an overly broad platform request obtaining data the user has not released. The second is access paralysis — a system so locked down that a genuine responder cannot reach a person in danger because the consent mechanism cannot be satisfied under degraded conditions. Blue Force Bridge resolves both by treating the authorization mechanism itself as a safety variable: when the user is able to authorize, they do so through a configured method of their choosing. When they are not, predefined governance rules — established when the user was not under duress — govern the response. The system does not guess. It executes a decision the user already made, under conditions the user anticipated, with a record that documents everything.

### Tactical Audio Bridge: Live Responder Communication Without Wiretap Risk

The Tactical Audio Bridge is the most sensitive functional component of the Blue Force Bridge architecture, and the most carefully bounded. Live voice communication between a responder and a person in danger is one of the highest-value safety interventions available to civilian infrastructure. A trained responder can provide breathing instructions, triage guidance, physical stabilization direction, and psychological grounding to a person who would otherwise be alone in a dangerous situation. The Tactical Audio Bridge is designed to enable that communication without converting it into a persistent surveillance artifact.

The architectural separation that makes this possible is key-material isolation. The live session is not an extension of the NICOLE evidence vault. It uses a separate ephemeral channel with separately derived key material, a distinct session identity, and a bounded lifetime tied to the specific access grant. The conversation does not automatically become part of any evidence record. It becomes part of a forensic record only if two independent conditions are met: configured recording authority has been granted, and that grant has been explicitly invoked and logged as a governance event. This is not the way most safety products handle audio. Most products that include live audio treat the session as part of a continuous monitoring record, or treat recording as the default state and require active opt-out. LAKANA reverses both assumptions. The default state is ephemeral, isolated, and non-recording. Recording requires a separate governance event with its own audit entry.

This design matters for reasons that extend beyond privacy. In post-incident proceedings — litigation, insurance disputes, internal use-of-force reviews, or regulatory investigations — the chain of custody and provenance of audio records is heavily scrutinized. A recording made without clear legal authority, without proper disclosure, or without a documented basis is not only potentially inadmissible; it is a liability for the institutions that collected it. LAKANA's architecture ensures that any audio that enters the evidence record does so through a governed, documented path that can withstand adversarial review. And audio that does not need to be a permanent record — the live coaching session between a trapped worker and a rescue coordinator, the check-in call between a rideshare driver and a safety contact — remains ephemeral, leaving no retrievable artifact that can be subpoenaed, leaked, or weaponized.

The session announcement mechanism addresses another scenario that is rarely discussed in safety system design: what happens when the device is in the hands of a hostile party and a responder attempts to initiate contact? Without announcement, an audio session could be opened silently, allowing a third party monitoring the responder channel to hear ambient audio that reveals location, occupancy, or the presence of a hidden user. LAKANA's announcement layer — the audible tone, voice prompt, or configured accessibility cue that notifies the user a request is incoming — is not just a user-experience feature. It is a security primitive. It prevents silent channel establishment. It gives the user, or any person in the environment, the ability to know that a communication attempt is in progress. And it creates a logged timestamp for the exact moment notification was delivered, which matters if the access is later contested.

**Coercion Resistance: When the Safety System Is the Attack Surface**

One of the most architecturally challenging scenarios in civilian safety is not the absence of a responder — it is the presence of a threat who understands that the safety system exists and attempts to weaponize it. An abusive partner who demands access to a location-sharing feature. A coercive employer who requires workers to grant monitoring permissions. A hostile passenger who attempts to force a rideshare driver to cancel an active safety protocol or reveal their current status. In each of these scenarios, the conventional safety product fails in a specific and predictable way: the threat has physical access to the device and social leverage over the user, which is sufficient to bypass most consumer consent mechanisms.

Blue Force Bridge is designed with this failure mode as a first-class design constraint, not an afterthought. The consent-airlock sequence does not recognize social pressure, physical device possession, or forced screen interaction as valid authorization. The system requires authorization through a configured method — a voice keyword, a gesture, a biometric trigger — that the threat cannot easily replicate or force without creating evidence of coercion. When the system detects unsafe-release conditions — a pattern of interaction consistent with duress, a rapid succession of canceled requests, input inconsistent with the user's established baseline — it narrows permitted actions rather than expanding them. It does not surface a prompt that asks the user to confirm under pressure. It holds, and it logs the hold.

This design also addresses a second coercion vector: the platform-level attack. An institution, a law enforcement agency without valid process, or a commercial party with access to a data broker relationship cannot simply request a Blue Force Bridge session and receive access. The bridge requires a verified responder identity, a declared safety function type, and a governance chain that traces back to a user-authorized configuration. There is no administrative portal that allows a privileged caller to bypass that chain. There is no law enforcement back door in the architecture. If

valid legal process is presented, the governed evidence custody path through NICOLE is the appropriate mechanism — not an emergency override of the consent airlock. LAKANA is not designed to resist lawful process. It is designed to ensure that any access, lawful or emergency, travels through a documented, auditable path rather than a silent override.

### Post-Incident Accountability and Chain of Custody

Blue Force Bridge's audit architecture is not a compliance feature. It is a safety feature. The ability to reconstruct exactly what access was granted, by whom, under what declared authority, at what time, and with what scope is valuable not only to the user but to every institutional partner that interacts with LAKANA data. An emergency management program that participated in a LAKANA-assisted incident response can produce a complete access log that documents their responder's actions from request to closure. An insurance carrier reviewing a rideshare incident can see a timestamped record of what location data was released, when, and under what governance rule — not a reconstructed account assembled after the fact from multiple systems. A legal team reviewing a field-work injury can see whether the user was conscious at the time of the access grant, what method they used to authorize it, and whether any coercion indicators were logged.

This auditability has direct commercial value. Institutions that adopt safety systems are increasingly exposed to liability for how those systems collect and handle sensitive data under emergency conditions. A system that produces a complete, tamper-evident access log from every Blue Force Bridge session is a materially different compliance posture than a system that produces server-side logs that can be modified, deleted, or simply never created if the network was unavailable. LAKANA's audit chain is local-first, meaning it does not depend on a cloud connection to produce a record. If the network was unavailable during an incident, the local state machine still recorded what happened. The audit chain is part of the safety substrate, not an optional reporting layer.

### Practical Adoption: Why Bounded Access Serves Responders Better

Blue Force Bridge creates a concrete adoption pathway for emergency management partners, firefighter academies, campus safety programs, field-service team supervisors, and rideshare safety pilots because its value proposition is immediately legible in operational terms. Consider the scenarios side by side. In a structural collapse, a rescue coordinator needs to reach a trapped worker. Without Blue Force Bridge, the options are either no contact at all, or access to the worker's entire communication history, location trail, and device state — a data surface that creates chain-of-custody problems and organizational liability. With Blue Force Bridge, the coordinator makes a scoped request for a live audio session and last-known location. The request is announced to the worker, authorized through a preconfigured emergency method or an emergency override, and the session begins — documented, bounded, and terminable at defined closure conditions.

In a rideshare incident, the scenario is different but the architecture is the same. A driver who has activated a safety protocol needs a connection to a configured safety contact or dispatcher — not a connection that exposes every prior trip, every route deviation, every passenger rating, and every home address associated with their account. The bridge provides the former without the latter. In a field-work accident in a basement or a remote utility corridor, a supervisor needs breathing status, scene instructions, and a transport confirmation — not a full behavioral profile of the worker. In a domestic violence scenario where a survivor has pre-configured specific trusted contacts and emergency release conditions, Blue Force Bridge is the mechanism that ensures only those contacts, under only those conditions, can access safety-relevant information — regardless of whether the

abuser is present, regardless of whether the device has been physically compromised, and regardless of whether social pressure is being applied at the moment of the incident.

The operational insight that makes Blue Force Bridge genuinely useful rather than merely architecturally correct is this: legitimate responders do not need unlimited access to do their jobs well. They need the right information at the right time under documented conditions. Unlimited access creates liability, evidentiary risk, and institutional exposure that professional responders and their organizations are trained to avoid. Bounded, logged, purpose-scoped access creates exactly the documentation and chain of custody that post-incident review demands. LAKANA is not limiting responders by building the consent airlock. It is giving responders a tool they can actually use in professional practice without creating the organizational and legal exposure that open-access systems impose.

# 6. Data Lifecycle, Context, and Coercion Resistance

## 6.1 Context Lanes and Temporal State

LAKANA's context lanes do not profile users. They prevent unsafe interpretation when the same signal means different things in different settings. A sudden stop on a highway, in a parking lot, inside a building, and during a storm-damaged route should not be treated as the same event. Context bounds decisions. It does not expand behavioral tracking.

At the architectural level, context enters as derived safety state: geographic sector, mesh sector, responder sector, last-safe sector, RF quiet map, threat sector, temporal ordering, and event-history plausibility. These lanes must be validated with field hardware before deployment claims are made. In the current evidence posture, they support architecture-level reasoning and pilot planning.

## 6.2 Coercion Resistance and Unsafe Release Prevention

Safety systems fail morally when they assume every user interaction is free and safe. A person may be forced to unlock a device, share a location, cancel an alert, or appear calm. LAKANA treats duress, ambiguity, and unsafe disclosure as first-class safety conditions. The system is designed to avoid exposing the user merely because someone nearby demands it.

The architecture uses several principles: keep raw data local by default, separate ordinary user interaction from evidence custody, permit silent evidence preservation when overt behavior could increase danger, make all release purpose-bound, and record coercion-sensitive state without creating a new stalking surface. The design goal is not to outsmart every attacker. The design goal is to avoid giving attackers a predictable surveillance surface.

## 6.3 Data Lifecycle: From Signal to Governed Retention

| Incident Scenario | Architectural Outcome |
|---|---|
| **Routine non-incident** | Minimal persistent residue. No permanent map feed. No audio retention created merely because a microphone-capable feature exists. |

| | |
|---|---|
| **Legitimate incident** | Evidence sufficient to reconstruct the safety event, hash-sealed and release-governed. Nothing more than what the governed release permits. |
| **Contested or coerced incident** | The fact of coercion risk is preserved without providing the coercive party a clean surveillance interface. Unsafe-release conditions narrow behavior. |
| **Failed transport attempt** | The failure is recorded as state, not erased. The audit story is complete regardless of network outcome. |
| **Responder access request** | Access is rescue-scoped, logged, and isolated from general user history. Denied access is also recorded as state. |

**WHAT THIS MEANS IN PRACTICE**

No permanent map feed is required for ordinary safety posture.

No raw audio retention is created merely because a microphone-capable feature exists.

No responder receives general-purpose user history as the default access mode.

No failed transport attempt disappears from the audit story.

No external dashboard becomes the source of truth unless a governed release path created that authority.

# 7. Companion Technical Preprint: Evidence Scope and Boundaries

The companion technical preprint — LAKANA SOS v71 Final Core Connected Empirical Monte Carlo Run — is archived at DOI 10.5281/zenodo.19956214 and provides the quantitative evidence base for this white paper. This white paper interprets that evidence at a non-reconstructive level. Full implementation details, protected thresholds, anti-abuse logic, and chain-of-custody internals are withheld for IP, safety, and misuse-prevention reasons.

**COMPANION PREPRINT CITATION**

**Fails, MarTaize K.** *LAKANA SOS v71: A Local-First, Fail-Closed Safety Operating System for Civilian Protection Under Degraded Infrastructure.* Technical preprint. LAKANA Sovereign Systems, 2026. DOI: 10.5281/zenodo.19956214

**White paper:** *Fails, MarTaize K. LAKANA Sovereign Safety System: Privacy-Preserving Civilian Safety Infrastructure for Protection Without Surveillance.* SSRN Working Paper v1.1, May 2026. Companion preprint DOI: 10.5281/zenodo.19956214.

## What the Simulation Establishes

The v71 run executed 250,000 full trials and 500,000 adversarial-stress trials under empirically calibrated parameters. The headline results, bounded by the stated assumptions, are as follows.

| Metric | LAKANA + CivOS | Industry Comparator |
|---|---|---|
| **Simulated safe-state survivability** | Higher than modeled | Modeled centralized/cloud, |

| | | |
|---|---|---|
| | comparator under specified assumptions — see preprint for precise values | single-uplink architecture |
| **Empirically calibrated transport success** | 0.939 | 0.674 |
| **Sample size** | 250,000 full trials + 500,000 adversarial-stress trials | Same population, same parameters |
| **Evidence classification** | Architecture-level simulation evidence | Architecture-level simulation evidence |
| **Field validation status** | Required — not yet completed | Required — not yet completed |

**EVIDENTIARY BOUNDARY**

The v71 results are architecture-level simulation evidence under specified modeled assumptions. They are not field validation, procurement certification, emergency-service approval, or a guarantee of real-world performance. They constitute a disciplined basis for advisor review, grant planning, pilot design, hardware-in-the-loop testing, and field benchmarking. The next evidence tier requires institutional partners.

## 8. Competitive Difference

The comparison below uses public architectural categories rather than making specific liability claims against any individual company. The point is structural: when a business model or technical architecture depends on continuous location extraction, the safety claim is architecturally compromised regardless of intent, policy promises, or marketing framing. The v71 modeled comparator represents a centralized/cloud, OS-intact, single-uplink architectural family.

| Dimension | Legacy Location / Telematics Model | LAKANA SOS Model |
|---|---|---|
| **Primary Posture** | Continuous monitoring, cloud visibility, or provider custody as normal operating assumptions. | Local protection as the default. Release is evidence-bound, purpose-bound, and governed. |
| **Business Incentive** | Data supports subscriptions, ads, analytics, broker sales, insurance scoring, or partner monetization. | Protection is the product. User safety data is not the product. |
| **Decision Logic** | Probabilistic inference or telematics scoring may be opaque to users and reviewers. | Deterministic bounded envelopes. Permitted actions are explainable and auditable. |
| **Infrastructure Failure** | A failed uplink can collapse alerting, logging, and escalation together. | Transport success means a packet found at least one permitted route. Failed attempts are recorded. |
| **Evidence Custody** | Logs may be server-side, mutable, incomplete, or provider-dependent. | NICOLE creates separated, hash-sealed, consent-aligned records under governed release. |
| **Coercion Risk** | Always-on location tools can become | Duress and unsafe-release conditions |

| | | |
|---|---|---|
| | control surfaces for abusers or coercive institutions. | narrow system behavior as first-class safety states. |
| **Responder Access** | May expose too much, too early, or without a clear evidence boundary. | Blue Force Bridge: rescue-scoped, consent-gated, isolated, and audit-governed. |
| **Transport Evidence (v71)** | Modeled comparator: 0.674 transport success under calibrated conditions. | LAKANA + CivOS: 0.939 transport success. Field validation required before deployment claims. |

# 9. What LAKANA SOS Can Do for Communities

LAKANA SOS is not a single consumer scenario. It is a civilian safety infrastructure layer for people whose risk is amplified by isolation, movement, infrastructure fragility, or contested evidence. Initial deployments should be controlled, local, and measurable, not mass consumer launches.

| Community | The Problem | LAKANA Safety Value |
|---|---|---|
| **Gig Workers & Rideshare Drivers** | Solo work, often at night, with strangers, in low-connectivity locations. Incident evidence is frequently disputed. | Consent-bound safety state, last-safe context, incident evidence custody, degraded transport, and responder bridge scoped to defined events. |
| **HVAC, Field-Service & Construction** | Workers enter basements, roofs, utility rooms, and storm-damaged structures where connectivity and visibility may be poor. | Local-first safety state, transport fallback, audit records, and team-visible incident packets without continuous worker surveillance. |
| **Emergency Management Drills** | Communities need last-mile communications and evidence workflows tested before disasters, not discovered during them. | Tabletop exercises validate degraded transport, audit chain, consent boundaries, and operator workflow in controlled conditions. |
| **Schools, Youth Programs & Athletics** | Parents and institutions need safety visibility without creating a permanent child-location marketplace. | Purpose-bound release, evidence minimization, strict non-surveillance assumptions, and consent-gated escalation. |
| **Domestic Violence & Stalking-Sensitive Populations** | Always-on tracking becomes a weapon when safety data reaches the wrong hands. | Coercion-aware modes, unsafe-release prevention, silent evidence preservation, and architecture that does not require public exposure. |
| **Community Resilience Partners** | Nonprofits and local responders need safety tools that support resilience without creating centralized behavioral control. | Local-first packets, partner dashboards under access control, and staged validation with institutional advisors before any deployment claims. |

# 10. Commercialization and Validation Ladder

The commercialization sequence must place institutional advisors before pilots. LAKANA currently has founder-built architecture, simulation evidence at DOI 10.5281/zenodo.19956214, prototype direction, and a provisional IP posture. The next evidence tier requires external judgment: advisors, study design, consent review, partner workflow review, and measurable validation endpoints. A small pilot without this advisory layer risks collecting data without producing credible evidence.

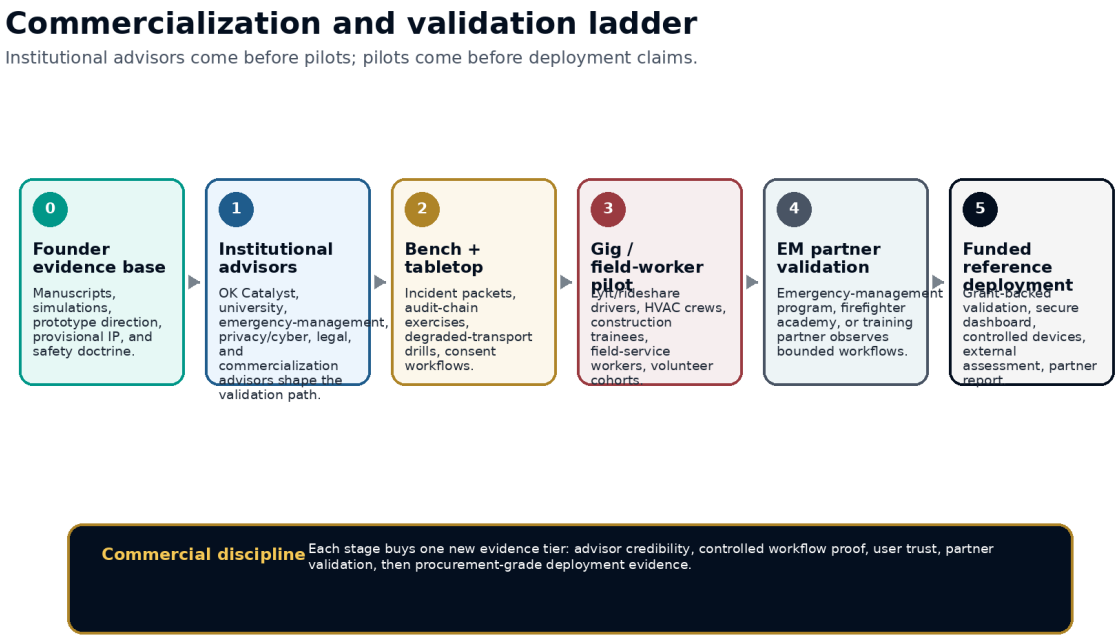


*Figure 4. Validation ladder: institutional advisors shape the path before pilots; pilots produce evidence before deployment claims.*

| | Phase | What Happens | Evidence Created |
|---|---|---|---|
| 0 | **Founder Evidence Base** | Manuscripts, simulation preprint (DOI 10.5281/zenodo.19956214), prototype direction, provisional IP, and safety doctrine. | Architecture rationale, simulation-level safety results, prior-art differentiation. |
| 1 | **Institutional Advisors** | OK Catalyst / SBIR triage, university, emergency management, privacy/cyber, legal, and commercialization advisors engaged. | Disciplined guidance before pilots. Prevention of overclaiming. Agency-fit map. |
| 2 | **Bench + Tabletop** | Incident packets, audit-chain exercises, degraded-transport tests, consent workflow evaluation. | Controlled workflow proof, measurable endpoints, partner-reviewed after-action reports. |
| 3 | **Gig / Field-Worker Pilot** | Rideshare drivers, HVAC crews, construction trainees, field-service workers, volunteer cohorts. Narrow, real, measurable conditions. | Validated pilot dataset, transport reliability data, consent workflow evidence, false-alert burden measurement. |
| 4 | **EM Partner Validation** | Emergency management program or academy. Tabletop/drill host, workflow | Practical operational fit, external assessment, public-sector reference |

| | | | |
|---|---|---|---|
| | | feedback, responder boundary review. | package. |
| 5 | **Funded Reference Deployment** | Phase II hardware integration, secure dashboards, controlled devices, external assessment, partner report. | Procurement-ready evidence, hardened prototype, deployment architecture, risk-control documentation. |

The first pilot lane should be narrow and real: gig and field-worker safety. Rideshare drivers, HVAC technicians, field-service crews, construction trainees, and volunteer cohorts create measurable safety conditions without claiming emergency-system replacement. LAKANA can measure transport reliability, consent workflow fit, false-alert burden, user trust, evidence-chain completeness, and usability under stress. The emergency management lane begins as tabletop and drill observation, not live public deployment.

## 11. Institutional Fit and Grant Alignment

LAKANA's strongest funding posture is a resilience, privacy, evidence-custody, and degraded-communications validation program. It aligns with multiple non-dilutive lanes because it addresses community protection, infrastructure fragility, trustworthy systems, emergency communications, and accountable safety workflows in a single integrated architecture.

| Funding Lane | LAKANA-Aligned Framing | What the Partner Helps Prove |
|---|---|---|
| **OK Catalyst / SBIR-STTR** | Convert founder-built evidence into agency-fit technical aims, budget narrative, milestones, and reviewer-ready claims. | Proposal discipline, advisor network, agency sequencing, Phase I readiness. |
| **DHS S&T / SBIR** | Prototype feasibility for local-first safety infrastructure under degraded, adversarial, or low-connectivity conditions. | Technical feasibility, measurable endpoints, privacy and cyber guardrails, prototype hardening. |
| **FEMA BRIC / HMA** | Partner-led resilience demonstration for last-mile safety communication, evidence continuity, and community risk reduction. | Eligibility fit through a local or state partner, mitigation framing, non-replacement boundaries. |
| **CISA / Emergency Communications** | Degraded-communications and interoperability-adjacent validation without claiming replacement of architecturally rated radio or 911. | Communication workflow relevance, tabletop validation, security review. |
| **University Research Partner** | Independent study design, statistics, human factors, cybersecurity, privacy review, and STTR-capable validation. | Credibility that a solo founder cannot ethically claim alone. Potential STTR eligibility. |
| **Emergency Management Program** | Controlled drills, responder workflow assessment, training-context feedback, after-action reports. | Practical operational fit without premature deployment claims. |
| **Legal / Privacy Advisors** | Consent architecture, data governance, retention, minor-safety posture, and abuse-prevention review. | Protection against repeating the surveillance-capitalism failure mode LAKANA is designed to prevent. |

**ADVISORY-FIRST POSTURE**

LAKANA cannot ethically claim independent validation, emergency-management suitability, privacy certification, or public-sector readiness alone. The correct posture is founder-built architecture plus external advisors, then controlled pilots, then externally reviewed evidence. This sequencing is what makes the evidence credible to institutional funders, not just internally coherent.

## 12. Limitations and Non-Claims

Credibility requires stating precisely where the evidence stops. The following limitations are stated explicitly, not placed in footnotes.

| Claim Category | Accurate Characterization |
|---|---|
| **v71 Simulation Results** | Architecture-level simulation evidence under modeled assumptions. Not field validation. Supports pilot planning and non-dilutive funding pursuit; does not support deployment certification. |
| **Industry Comparator** | Represents a centralized/cloud, OS-intact, single-uplink architectural family. Does not represent every competitive configuration. |
| **Transport Success Result** | A pilot target, not a field guarantee of connectivity in any specific environment. |
| **TSARO Accuracy** | Meaningful but not perfect. Validation must improve threat-conditioned sensitivity while reducing false escalation and alert fatigue. |
| **NICOLE Integrity Behavior** | Simulation behavior does not equal independent cryptographic certification, privacy audit, or legal admissibility ruling. |
| **GUIDE / Temporal Inputs** | v71 evidence includes proxy-derived components. Hardware-captured field validation is required. |
| **Blue Force Bridge / Audio** | Must remain consent-gated, rescue-scoped, logged, and externally reviewed before any public-sector deployment. |
| **System Scope** | Not a medical device, law-enforcement tool, or replacement for 911, emergency managers, or incident command. |
| **Certifications** | No proposal should claim FEMA approval, NIMS/ICS certification, field-validated life-safety guarantee, or independent privacy certification unless those validations are completed separately. |

## 13. Conclusion: The Architecture the World Needs Now

The world does not need another application that maps people more aggressively and calls it safety. What is needed is safety infrastructure that becomes more disciplined as risk increases, treats evidence as a safety object, governs before it collects, and does not require users to surrender their identity, routine, and vulnerability as the price of protection.

The reason this matters urgently is the accelerating capability of artificial intelligence. AI makes inference cheaper, faster, and more invasive with each passing year. The same telemetry that once

showed a dot on a map can now deduce routine, vulnerability, health, relationships, religious practice, labor exposure, domestic risk, and financial standing. A safety company that collects everything and promises to behave is structurally insufficient. Promises are not architecture. The architecture itself must prevent safety from becoming surveillance.

The documented pattern — Life360 supplying data brokers, Arity building behavioral scoring from embedded mobile SDKs, the FTC pursuing four separate enforcement actions against location-data commercial chains in two years, a Texas AG suit alleging insurance-rate consequences for 45 million Americans — is not a series of isolated incidents. It is the predictable output of an architecture that begins with maximum data collection and attempts to govern it afterward. LAKANA begins at the opposite end: govern first, collect only what governed release permits, preserve evidence without creating surveillance, and make every safety action explainable and auditable.

**THE IMMEDIATE ASK**

LAKANA SOS is positioned for the next step: advisor-backed proposal strategy, disciplined non-dilutive funding, controlled validation, and public-sector and commercial partner review.

The immediate ask is not blind trust in founder claims. It is institutional help to prove the architecture correctly: narrowly, rigorously, ethically, and without compromising the people it is meant to protect.

The companion technical preprint (DOI: 10.5281/zenodo.19956214) provides the quantitative foundation. This white paper provides the rationale, the architecture, the evidence boundaries, and the path forward. Together they constitute a foundation for the next evidence tier — one that no founder can ethically build alone.

## Data and Code Availability

The companion technical preprint, LAKANA SOS v71 Final Core Connected Empirical Monte Carlo Run, is publicly archived at DOI 10.5281/zenodo.19956214. The archived preprint discloses aggregated simulation results, figure-level summaries, evidence-tier descriptions, parameter ranges, claim boundaries, and a public-safe interpretation of the architecture at a non-reconstructive level of detail.

Full implementation source code, protected threshold schedules, anti-abuse logic, private chain-of-custody protocol internals, and reconstructive engine details are withheld for intellectual-property, operational security, and misuse-prevention reasons. These materials are not required to evaluate, reproduce, or audit the claims stated in this white paper or the companion preprint at the disclosed evidence tier.

Controlled-access review of restricted technical materials may be pursued under appropriate confidentiality, legal, and safety terms by contacting the author directly. Institutional partners pursuing pilot validation, grant co-investigation, or procurement-readiness review may request a scoped technical disclosure under a mutually executed non-disclosure and responsible-use agreement.

## AI-Assisted Drafting Disclosure

The author used AI-assisted drafting, editing, formatting, and publication-preparation support in the development of this white paper. The author reviewed, directed, and approved all final content and remains solely responsible for all claims, citations, stated evidence boundaries, limitations, and submission decisions presented in this document. No AI system is listed as an author or co-author. This disclosure is provided in compliance with SSRN submission requirements and the author's commitment to transparent research practice.

## Funding Statement and Competing Interests

This work is founder-developed and privately bootstrapped by the author. LAKANA Sovereign Systems has received no external funding, grants, investment, or sponsored research support for the work described in this white paper at the time of publication. No third-party funding source has influenced the technical claims, evidence boundaries, architectural design decisions, or validation posture described herein.

The author is the founder and principal architect of LAKANA Sovereign Systems and holds a direct financial interest in the commercialization of the systems described in this document. This relationship is disclosed here in full. The author has no other relevant financial interests, consulting relationships, or affiliations that would constitute a competing interest with respect to the claims made in this white paper. LAKANA Sovereign Systems: www.lakana.systems.

## References

1. Keegan, Jon. "The Popular Family Safety App Life360 Is Selling Precise Location Data on Its Tens of Millions of Users." The Markup, December 6, 2021. https://themarkup.org/privacy/2021/12/06/the-popular-family-safety-app-life360-is-selling-precise-location-data-on-its-tens-of-millions-of-users (accessed May 2026).
2. Keegan, Jon. "Life360 Says It Will Stop Selling Precise Location Data." The Markup, January 27, 2022. https://themarkup.org/privacy/2022/01/27/life360-says-it-will-stop-selling-precise-location-data (accessed May 2026).
3. E.S. v. Life360 Inc., Complaint, No. 4:23-cv-00168, U.S. District Court for the Northern District of California, January 12, 2023. Filed January 12, 2023; voluntarily dismissed with prejudice. PACER CourtLink, N.D. Cal.
4. ClassAction.org. "Life360 Secretly Sells Users' Geolocation Data to Third Parties, Class Action Claims." https://www.classaction.org/news/life360-secretly-sells-users-geolocation-data-to-third-parties-class-action-claims (accessed May 2026; status updated February 9, 2026).
5. Texas Attorney General. "Attorney General Ken Paxton Sues Allstate and Arity for Unlawfully Collecting, Using, and Selling Over 45 Million Americans' Data." January 13, 2025. https://www.texasattorneygeneral.gov/news/releases/attorney-general-ken-paxton-sues-allstate-and-arity-unlawfully-collecting-using-and-selling-over (accessed May 2026).
6. Reuters. "Texas sues Allstate over data collection from cellphones," January 13, 2025; Reuters. "Allstate must face privacy lawsuit over cellphone tracking of drivers," March 4, 2026. Reuters.com (accessed May 2026). Allstate has denied the allegations and asserted consent-based defenses.
7. Federal Trade Commission. "FTC Order Prohibits Data Broker X-Mode Social and Outlogic from Selling Sensitive Location Data." January 9, 2024; final order April 2024. https://www.ftc.gov/news-events/news/press-releases/2024/01/ftc-order-prohibits-data-broker-x-mode-social-outlogic-selling-sensitive-location-data (accessed May 2026).
8. Federal Trade Commission. FTC v. Kochava, Inc., No. 2:22-cv-00349 (D. Idaho). Complaint filed August 29, 2022. Case materials available at https://www.ftc.gov/legal-library/browse/cases-proceedings/2023186-kochava (accessed May 2026).

9. Federal Trade Commission. "FTC Finalizes Order Prohibiting Gravy Analytics, Venntel from Selling Sensitive Location Data." January 14, 2025. https://www.ftc.gov/news-events/news/press-releases/2025/01/ftc-finalizes-order-prohibiting-gravy-analytics-venntel-selling-sensitive-location-data (accessed May 2026).
10. Federal Trade Commission. "FTC Takes Action Against Mobilewalla for Collecting and Selling Sensitive Location Data." December 3, 2024. https://www.ftc.gov/news-events/news/press-releases/2024/12/ftc-takes-action-against-mobilewalla-collecting-selling-sensitive-location-data (accessed May 2026).
11. Federal Emergency Management Agency. National Incident Management System, Third Edition. October 2017. https://www.fema.gov/sites/default/files/2020-07/fema_nims_doctrine-2017.pdf (accessed May 2026).
12. Cybersecurity and Infrastructure Security Agency. National Emergency Communications Plan (NECP), 2019. https://www.cisa.gov/necp. SAFECOM Interoperability Continuum, 2021. https://www.cisa.gov/safecom (accessed May 2026).
13. National Institute of Standards and Technology. Cybersecurity Framework 2.0. NIST CSWP 29, February 26, 2024. DOI: 10.6028/NIST.CSWP.29. https://doi.org/10.6028/NIST.CSWP.29.
14. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1, January 2023. DOI: 10.6028/NIST.AI.100-1. https://doi.org/10.6028/NIST.AI.100-1.
15. National Institute of Standards and Technology. Engineering Trustworthy Secure Systems. NIST SP 800-160 Vol. 1 Rev. 1, November 2022. DOI: 10.6028/NIST.SP.800-160v1r1. https://doi.org/10.6028/NIST.SP.800-160v1r1.
16. Fails, MarTaize K. LAKANA SOS v71: A Local-First, Fail-Closed Safety Operating System for Civilian Protection Under Degraded Infrastructure. Technical preprint. LAKANA Sovereign Systems, 2026. Zenodo. DOI: 10.5281/zenodo.19956214. https://doi.org/10.5281/zenodo.19956214.
17. Fails, MarTaize K. LAKANA Sovereign Safety System: Privacy-Preserving Civilian Safety Infrastructure for Protection Without Surveillance. SSRN Working Paper v1.1, May 2026. Companion technical preprint DOI: 10.5281/zenodo.19956214.
18. Blue Force Bridge V2.5 Technical Specification. LAKANA Sovereign Systems, 2026. Internal technical specification. Available to institutional partners under NDA.