

Naval Base San Diego

National Security Exposure Assessment


Joshua Michael

Email: josh@nexanet.ai



© 2026 Joshua Michael · [CC BY 4.0](https://creativecommons.org/licenses/by/4.0/). Use with credit.

September 2, 2026



Executive Assessment

Naval Base San Diego is the Navy's main Pacific surface-fleet homeport — ships, crews, and the Southwest Regional Maintenance Center that keeps them running. A crew, maintainers, and contractors starting to move together around one ship could mark a deployment or a major repair milestone.

This study asks a simple question: if someone drove from a nearby home to a gate and back, would they pass a Flock camera? The numbers are a model of possible drives, not a count of real people.

At 20 miles, about 67.0% of those modeled trips pass a camera. Another 15.7% might pass one, 2.2% avoid them, and 15.1% could not be determined. Even at 5 miles, about 74.2% of those trips still pass a camera. This is a model of routes, not proof that anyone was watched.

HEADLINE RESULT · MODELED TRIPS TO A GATE

20 MILES OUT



■ **67.0%**

Passes a camera

A confirmed camera sits on every plausible road to a gate.

■ **15.7%**

Might pass

The inventory's rotationAngle is not a true facing: among 72,971 named in-service cameras it matches the road only about 10% of the time (median error 90°) because it stores housing orientation, not travel. The model therefore does not use it to confirm a hit. Named NB/SB/EB/WB cameras take direction from the road, with the name used only to pick which way they watch. Cameras with no compass name are too noisy to confirm, so a nearby pass stays Might pass. A confirmed hit still needs the vehicle on the same corridor, traveling in that monitored direction within 35°.

■ **2.2%**

Avoids all cameras

No confirmed camera on any plausible road.

■ **15.1%**

Couldn't determine

Extra driveway or side-road options could not be fully routed, so those trips stay unresolved instead of being guessed.

5 MILES OUT



■ **74.2%** Passes ■ **7.8%** Might pass ■ **1.6%** Avoids ■ **16.4%** Couldn't determine

Couldn't determine is 16.4% here because extra driveway or side-road options could not be fully routed, so those trips stay unresolved instead of being guessed.

Installation Context

Naval Base San Diego is the Navy's principal Pacific Fleet surface-force installation and supports homeported ships, tenant commands, regional headquarters, medical functions, family housing, and a large military-civilian-contractor population [E104, E105]. The Southwest Regional Maintenance Center is publicly located at Naval Base San Diego and contributes to the Navy mission of maintaining ships and systems [E171, E172].

These missions create two separate movement risks. A homeported surface combatant preparing to deploy may generate a recurring crew, maintenance, logistics, security, and command-support signature. SWRMC and contractor personnel form a smaller technical network whose changing attendance can reveal ship-specific repair, modernization, inspection, or readiness activity.

The study did not identify a ship, sailor, SWRMC employee, contractor, maintenance issue, deployment, or household. The four reviewed gates are model inputs, and synthetic route outcomes cannot establish the actual access habits of any ship's company or tenant organization.

Installation-Specific Intelligence Scenarios

The study demonstrates Tier A deployment and modeled route exposure only. Both scenarios below require Tier B historical time-stamped observations and Tier C identity or organizational enrichment that were not obtained.

Scenario 1: Detecting a Homeported Surface Combatant's Deployment or Readiness Surge

Assume that a foreign service obtains historical plate observations around Naval Base San Diego and identifies a recurring cluster associated with one homeported surface combatant and its support organizations. The cluster may include crew members, maintainers, supply personnel, security, command staff, medical support, and contractors.

A deployment or major readiness event may produce an observable sequence: increased technical attendance, changed shift patterns, added logistics traffic, prolonged command activity, and the eventual disappearance of a large part of the crew cluster. A later coordinated return would provide an approximate mission-completion date.

The adversary might not know the ship's destination or tasking. It could still estimate when the vessel became unavailable at homeport, whether the event was routine training or a longer operation, and which support capabilities were involved. Correlation with public port reporting, satellite imagery, aviation activity, diplomatic events, or overseas crises could narrow the assessment.

The history also creates a recognition model. After observing workups and public deployments, the service can identify which movement changes are ordinary and which precede a real departure. The base's large population provides noise, but repeated ship-specific association can reduce it.

The GRU Unit 26165 campaign provides a direct precedent for this collection objective. Partner governments reported that Russian military intelligence targeted internet-connected cameras near military installations and logistics routes to monitor aid movement [E148]. The cameras did not need access to cargo documents; timing and physical movement were the intelligence.

The study did not identify a ship, crew, or deployment. It models route exposure only. The scenario requires longitudinal observations, ship attribution, and validation that the relevant population uses the modeled gates and public-road corridors.

Scenario 2: Mapping the Southwest Regional Maintenance Center and a Ship-Specific Contractor Team

SWRMC provides a more concentrated technical target than the general base population. Its public mission is tied to maintaining Navy ships and systems, and official reporting places the command at Naval Base San Diego [E171, E172]. Engineers, inspectors, artisans, technical representatives, suppliers, and contractors may recur around a particular ship or availability.

Historical plate co-occurrence could identify vehicles that repeatedly bridge SWRMC, the base, contractor facilities, parts suppliers, training locations, and other regional Navy sites. A sudden rise among one cluster may indicate an urgent casualty repair, combat-systems integration problem, inspection, modernization event, or return-to-service milestone.

Even when the technical issue remains classified, movement can reveal which companies and specialists are involved and how long the work persists. A foreign service could then focus cyber collection or human targeting on a small number of people likely to know the failure, limitation, or fielding timetable.

The OPM breach provides a direct precedent for state collection against personnel with sensitive access. GAO reported compromise of background-investigation data affecting 21.5 million people, and the FBI has attributed the theft to Chinese actors [E149, E150]. Such records could identify engineers or cleared contractors; ALPR history could identify the ship and current maintenance event with which they are associated.

The combination would allow retrospective analysis long after the original plate scans were collected. Once a ship problem became public, an adversary could query backward to reconstruct who supported the response and then watch those same specialists in future events.

The present research did not obtain SWRMC rosters, contractor lists, maintenance records, registration data, or historical plate reads. It cannot identify a ship-specific problem. The scenario explains why a technical-support workforce at a major fleet base should be treated as part of the operational attack surface.

Combined Significance

Scenario 1 reveals when a combatant is preparing to leave or return. Scenario 2 reveals the technical and industrial network sustaining that combatant and may expose why its readiness changed. Together they create a richer fleet-intelligence picture than ship tracking alone.

Primary Reachability Findings

Modeled route-coverage estimate, reported for both scopes with its uncertainty ranges.

Field	5-mile near-base	20-mile primary
Result label	V7 modeled estimate	V7 modeled estimate
Publication status	Modeled estimate (3 caveats)	Modeled estimate (3 caveats)
Headline channel	robust_lattice	robust_lattice
Primary H0 U/C/P/R	2.2%/87.5%/6.5%/3.8%	2.6%/75.3%/16.1%/5.9%
Robust H0/H1/H2 U/C/P/R	1.6%/74.2%/7.8%/16.4%	2.2%/67.0%/15.7%/15.1%
H0 validation	single-reviewer check	single-reviewer check

The complete modeled result record for both scopes - all sensitivity intervals, gate/direction/range sweeps, camera-cluster influence, modeling caveats, seeds, and SHA-256 hashes - is in the V7 Modeled Result Record under Evidence and Reproducibility below.

20-mile minus 5-mile weighted unobserved-rate difference: primary conditional H0 **+0.4** percentage points; separate robust H0/H1/H2 lattice **+0.6** percentage points. These channels answer different questions and are not averaged. [L41]

Legislative Significance

Naval Base San Diego warrants sensitive-site rules covering both fleet movement and maintenance networks. Legislators should require short non-hit retention, restricted historical and reverse searches, narrow sharing, vendor and subcontractor security controls, immutable audit logs, and immediate Navy notification after compromise. Public products should omit exact gate and corridor vulnerabilities.

Neither recommendation depends on a claim that a current operator acted improperly. The policy issue is whether a longitudinal commercial movement archive surrounding a sensitive national-security site should be retained, broadly shared, or exposed through vendor and subcontractor systems without controls proportionate to the intelligence value of the aggregate record.

Four Reported Vulnerabilities

Finding	Observed condition	Evidence boundary	Disclosure and remediation status	Source(s)
1	An active ArcGIS Default API key was hardcoded in a public development/demo asset and repeated key material was found across development assets.	The validation establishes key scope and private item grants; it does not prove that every granted item was read or altered by another party.	Reported 2025-06-11; remediation reported 2025-07-01.	L2, L5, L8, L20
2	FlockSafety development application and module-federation bundles were served before the correct authentication boundary.	The bundles expose application logic, routes, models, and permission concepts; frontend visibility does not prove anonymous access to every backend or production record.	Reported 2025-11-13 with follow-ups on 2025-11-14 and 2025-11-19; no closure or remediation detail appears in the disclosure record.	L2, L5, L11-L13
3	Planner/Beefeater endpoints issued ArcGIS and Google mapping credentials before login, including a production-scope path used to obtain the analyzed export.	Wire records prove anonymous HTTP 200 token responses and establish provenance; they do not prove every downstream service permitted every operation.	Reported 2025-11-13 with follow-ups on 2025-11-14 and 2025-11-19; no closure or remediation detail appears in the disclosure record.	L2, L14, L15, L37

Finding	Observed condition	Evidence boundary	Disclosure and remediation status	Source(s)
4	FlockNova application bundles were served before the correct access boundary and revealed search, pivot, fusion, and client-side permission assumptions.	The exposed code establishes application structure and data-source concepts; it does not establish that every modeled backend endpoint lacked authorization or that every modeled dataset was populated.	Reported 2025-12-05; no reply or remediation detail appears in the disclosure record.	L2, L4, L35

The findings are distinct but composable. Findings 2 and 4 exposed the vocabulary and architecture of search, alerting, mapping, and fusion workflows. Findings 1 and 3 exposed two different credential failures: a static production-scoped ArcGIS key and a pre-login credential-issuance path. Finding 3 supplied the provenance for the production map-point export used here. The combination supports a conclusion about systemic credential and environment-boundary weakness during the research period, while remaining narrower than a claim of universal platform compromise [L2, L8, L37].

The disclosure record also cannot establish a negative. Responsible reporting and a known remediation date for Finding 1 do not prove that no other person accessed the exposed surfaces, just as the absence of a returned closure notice for Findings 2 through 4 does not prove that they remain exploitable today. The reports preserve the observed condition, collection date, and known response status so that legal and technical reviewers can distinguish evidence from present-tense assumptions.

Methodology

Research and Data Basis

This assessment follows four separate disclosures involving credentials, development assets, environment boundaries, and a data-fusion application. A public development artifact contained an active ArcGIS credential; other development bundles exposed application structure before the expected access boundary; a Planner/Beefeater path issued production-scope mapping credentials before login; and FlockNova bundles exposed the structure of a broader fusion and search environment [L2, L4, L5, L8, L37]. None of those findings alone proves universal backend access. Together, they justified examining what sensitive geospatial information became available and what defensive questions it could support.

The production ArcGIS export captured on December 14, 2025 contains 335,701 records spanning multiple device types and statuses; it is not a count of simultaneously active cameras. The status fields record 214,073 `inService`, 83,427 `inPlanning`, 38,190 decommissioned, and 11 unknown records. The route study loads detection-eligible `inService` and `inPlanning` records and preserves a separate `inService`-only floor. The `active` integer is not the study filter [L14, L15].

Exposed application materials model plate and vehicle search, camera selection, hit history, alerts, people-search paths, map layers, CAD and AVL concepts, Raven/audio sensors, device access, and identity-adjacent fusion permissions [L4, L11-L13, L35]. Flock's patent describes a dynamic surveillance and object-query architecture, but a patent describes claimed architecture rather than proving deployment of each patented capability [E37]. The report therefore separates observed code and export fields from product descriptions and hypothetical intelligence use.

The force-protection concern arises because personnel assigned to Naval Base San Diego routinely use civilian roads outside a protected perimeter. National evidence indicates that many service members live in local communities and use civilian housing markets when government quarters are not provided [E44-E46]. Those national facts do not reveal where a person assigned to Naval Base San Diego lives; they establish only that off-installation movement is a legitimate research surface.

Longitudinal location information can expose associations and patterns that isolated observations do not. Supreme Court decisions concerning long-term GPS tracking and historical location records provide relevant policy and legal context without deciding the legality of this vendor, dataset, or study [E20, E21]. The constitutional and privacy implications belong in legal and policy review, not in unsupported declarations that every collection or search is unlawful.

My analytical judgment is that commercial growth and procurement convenience must not displace United States national-security, counterintelligence, force-protection, constitutional, and public-accountability obligations. That is a normative conclusion, not evidence that profit caused a particular breach or that a specific customer misused data.

Estimand and Synthetic Origins

The V7 primary estimand is the occupied-housing-weighted distribution across unobserved, confirmed exposure, possible exposure, and unresolved model under one predeclared, camera-blind H0 access-and-route policy. H0 defines the pinned origin, eligible unchanged saved gates, 5/20-road-mile frame, final per-leg scope, and primary conditional result. Every in-frame sampled weight remains in the four-class denominator; proven out-of-frame origins are excluded before sampling, while frame-unknown, unresolved-model, and artifact-error states remain explicit.

Each directed outbound and return leg must be no more than the declared radius. A qualifying round trip can therefore approach 10 miles in the 5-mile scope or 40 miles in the 20-mile scope. This is not a total-round-trip cap of 5 or 20 miles.

H1/H2 remain identity-keyed alternative access policies used only for the separate robustness lattice and fixed-H0-gate access sensitivity. H1/H2 cannot rescue a failed H0, set the primary status, change H0 frame membership, or contribute cameras to the H0 result. The two channels are co-located and never averaged [L41].

Reviewed 24/7 Gates

The paired study used Gate 1, Gate 2, Gate 3, and Gate 4 as the reviewed gate set. Only gates treated during setup as continuously open to authorized traffic are eligible. A part-time, visitor-only, commercial-only, or schedule-limited gate would create a modeled route that disappears outside its operating window, so including it as continuously available would overstate practical access. The gate decision concerns availability, not secrecy; coordinates and identifiers remain in the unredacted annexes.

The same gate set was used at both radii. The saved study state also contained 0 analyst-defined road blocks. Gates and blocks are inputs rather than outputs. A high or low unobserved rate does not validate the gate review, so an installation reviewer should read the exact saved identifiers and artifact warnings when judging whether the modeled access assumptions match local practice.

Cameras, Roads, and Routing

The V7 route and gate policy receives no camera coordinates, encounter results, or camera-derived costs. It uses the frozen routing and inventory assumptions bound into the completed cohort (`valhalla-3.6.3-auto-frozen-osm-camera-blind-directed-overlay-v6`). Outbound and return routes are legal directed requests and are evaluated independently against the declared per-leg scope.

Camera encounters remain geometry, corridor, bearing, lifecycle-status, and capture-model inferences rather than observed plate reads. The H0 point channel retains direction and spatial/grade sensitivity separately. Camera influence is computed by physical camera cluster so colocated compute/support records are not counted as independent roadside sensors. Planned-device and in-service-only status variants remain sensitivity evidence rather than claims of confirmed real-world operation [L41].

Why Five and Twenty Miles

DoD Manual 4165.63 permits housing-market boundaries based on a twenty-mile or one-hour peak-traffic commute, with military and public information and installation-command input supporting local adjustment [E49]. The study predeclared the distance limb as a uniform, transparent approximation of the broader off-installation housing and commuting area. It does not assert that most personnel at Naval Base San Diego live exactly twenty miles away, and it does not model one hour of peak traffic.

Five miles isolates the immediate area where travel converges toward a limited gate set. Twenty miles adds the broader housing-market stratum while retaining every near-base draw unchanged. Reporting both shows whether corridor concentration near the facility differs from the surrounding occupied-housing geography without changing the near-base sample.

Limitations and Non-Claims

This is a model over dated deployment, Census, and road snapshots. Occupied housing is a residential-distribution proxy, not an installation workforce roster; it does not weight origins by military assignment, employer, rank, duty status, shift, or measured traffic. National off-base housing evidence supports the scope choice but cannot identify where personnel assigned to Naval Base San Diego live [E44-E46].

Camera location and status fields can be incomplete, stale, planned, duplicated, moved, or differently interpreted. OSM roads and address points can omit valid travel or include obsolete topology. Gate availability and cantonment geometry depend on analyst review. The route engine models corridor blocking, accepted alternatives, and a geometric field-of-view capture; it does not model OCR probability, optical obscuration, weather, plate condition, device uptime, data retention, customer sharing, legal process, or whether a backend user queried a record.

A V7 unobserved H0 classification is a conditional modeled-policy result: under the frozen H0 access, gate, routing, inventory-status, and capture rules, the legal directed round trip had no qualifying modeled encounter. It is not a share of people, drivers, commuters, households, actual trips, or empirical camera reads. The separate robust lattice and access, gate-choice, direction, spatial/grade, range, deployment-status, sampling, and camera-cluster sensitivities must remain visible qualifications. Source omissions, device uptime, weather, plate visibility, temporary road conditions, and camera movement can change real-world collection.

Evidence and Reproducibility

The V7 drafting packet was built from the complete validated cohort for `naval_base_san_diego`. It binds cohort `reachability-v7-audit-92bd2e6a45c3a7c6`, pair `naval_base_san_diego-208108b789368d8b`, seed `2148787769`, index SHA-256

3d0755cb6fcc7dd165c6b9dd53ade8bfe5402773231904ea3807d1f17dc3b96b , five-mile summary SHA-256 7b35a3241fd3436e50a345dcfe6d09900496da0db8f547d020f79b96d6816c58 , and twenty-mile summary SHA-256 0ed2ff7ddbc73b03cacc460a01726a5f5b77f37b1009ffc64f3c7d1dcf909961 . [L41].

Both scopes retain the primary H0 and robust-lattice partitions, full frame accounting, per-policy outcomes and failures, directed route geometry, modeled camera encounters, split uncertainty intervals, gate-choice sensitivity, physical-camera-cluster influence, frozen source snapshot, and independent QA. All QA flags and warnings remain unsuppressed in the hash-bound V7 snapshot [L41].

The narrative is subordinate to those structured artifacts. If a sentence, table cell, or later export conflicts with a canonical summary or origin-level record, the validated structured record controls and the narrative must be corrected. Joshua Michael independently reviews each material claim, citation, inference, and numerical value as the accountable author.

V7 Modeled Result Record

Modeled route-coverage estimate [L41].

Field	5-mile near-base	20-mile primary
Result label	V7 modeled estimate	V7 modeled estimate
Publication status	Modeled estimate (3 caveats)	Modeled estimate (3 caveats)
V7 cohort	reachability-v7-audit-92bd2e6a45c3a7c6	reachability-v7-audit-92bd2e6a45c3a7c6
Pair ID	naval_base_san_diego-208108b789368d8b	naval_base_san_diego-208108b789368d8b
Pair seed	2148787769	2148787769
Generated at	2026-08-12T13:46:39.652260+00:00	2026-08-12T13:46:39.652260+00:00
Summary SHA-256	7b35a3241fd3436e50a345dcfe6d09900496da0db8f547d020f79b96d6816c58	0ed2ff7ddbc73b03cacc460a01726a5f5b77f37b1009ffc64f3c7d1dcf909961
Routing engine	valhalla-3.6.3-auto-frozen-osm-camera-blind-directed-overlay-v6	valhalla-3.6.3-auto-frozen-osm-camera-blind-directed-overlay-v6
Synthetic draws	1000	2000
Headline channel	robust_lattice	robust_lattice
Primary H0 U/C/P/R counts	22/875/65/38	49/1607/243/101
Primary H0 U/C/P/R	2.2%/87.5%/6.5%/3.8%	2.6%/75.3%/16.1%/5.9%
Robust H0/H1/H2 U/C/P/R counts	16/742/78/164	39/1400/248/313
Robust H0/H1/H2 U/C/P/R	1.6%/74.2%/7.8%/16.4%	2.2%/67.0%/15.7%/15.1%
Direction H0 interval	2.2% to 12.5%	2.6% to 24.7%

Field	5-mile near-base	20-mile primary
Fixed-gate access interval	1.6% to 18.9%	2.2% to 17.7%
Joint interval	1.6% to 25.8%	2.2% to 33.0%
Spatial H0 interval	1.9% to 5.1%	1.8% to 7.0%
Gate reselection diagnostic	fixed 1.7%; reselected 1.7%; absolute difference 0.0 pp	fixed 2.2%; reselected 2.2%; absolute difference 0.0 pp
H0 gate near-tie sensitivity	60s: 3.1% disagreement weight; 120s: 5.1% disagreement weight; 300s: 6.3% disagreement weight	60s: 1.4% disagreement weight; 120s: 3.1% disagreement weight; 300s: 8.9% disagreement weight
H0 range sensitivity U/C/P/R	75ft: 2.2%/87.3%/6.7%/3.8%; 150ft: 2.2%/87.5%/6.5%/3.8%; 250ft: 2.1%/87.7%/6.4%/3.8%	75ft: 2.7%/74.8%/16.6%/5.9%; 150ft: 2.6%/75.3%/16.1%/5.9%; 250ft: 2.4%/75.8%/15.9%/5.9%
H0 direction sensitivity U/C/P/R	25deg: 2.3%/87.4%/6.5%/3.8%; 30deg: 2.3%/87.5%/6.4%/3.8%; 35deg: 2.2%/87.5%/6.5%/3.8%; 40deg: 2.2%/87.5%/6.5%/3.8%; 45deg: 2.0%/87.5%/6.7%/3.8%; 60deg: 1.9%/87.5%/6.8%/3.8%	25deg: 3.1%/75.3%/15.7%/5.9%; 30deg: 2.7%/75.3%/16.0%/5.9%; 35deg: 2.6%/75.3%/16.1%/5.9%; 40deg: 2.6%/75.3%/16.1%/5.9%; 45deg: 2.6%/75.3%/16.2%/5.9%; 60deg: 2.1%/75.3%/16.7%/5.9%
Deployment-status sensitivity U/C/P/R	declared_policy: 2.2%/87.5%/6.5%/3.8%; in_service_only: 2.2%/87.5%/6.5%/3.8%	declared_policy: 2.6%/75.3%/16.1%/5.9%; in_service_only: 2.6%/75.3%/16.1%/5.9%
Physical-camera-cluster influence	maximum removal shift 3.2 pp; 0 material clusters	maximum removal shift 0.5 pp; 0 material clusters
H0 validation	single-reviewer check	single-reviewer check
Modeling caveats	1. corridor routing is a single-reviewer accuracy check (~89.5%; see the report's Appendix O)2. primary H0 unresolved weight is 3.800%3. fixed-gate access-policy interval spans 17.300 pp	1. corridor routing is a single-reviewer accuracy check (~89.5%; see the report's Appendix O)2. primary H0 unresolved weight is 5.934%3. fixed-gate access-policy interval spans 15.508 pp
QA and artifact warnings	None	None

20-mile minus 5-mile weighted unobserved-rate difference: primary conditional H0 **+0.4** percentage points; separate robust H0/H1/H2 lattice **+0.6** percentage points. These channels answer different questions and are not averaged. [L41]

Reachability Results

Five-Mile Near-Base Scope

Publication state: Modeled estimate (3 caveats).

The 5-mile primary H0 channel contains 1,000 draws on the full in-frame sampled denominator: 22 unobserved (2.2%), 875 confirmed exposure (87.5%), 65 possible exposure (6.5%), 38 unresolved model (3.8%). The separate robust H0/H1/H2 lattice reports 16 unobserved (1.6%), 742 confirmed exposure (74.2%), 78 possible exposure (7.8%), 164 unresolved model (16.4%). Neither channel drops or renormalizes unresolved-model weight.

H0 uncertainty intervals: direction-H0 2.2% to 12.5%; fixed-H0-gate access 1.6% to 18.9%; joint 1.6% to 25.8%; spatial/grade-H0 1.9% to 5.1%. These are distinct sensitivity results and are not averaged.

Candidate-frame accounting: {"all_saved_gates_evaluated": true, "certificate_artifact_path": "frame_classification.json", "certificate_artifact_sha256": "36907288f4cf7a4ccfabd54f0ee709028695812e693a2a1e915dcc25e6f092c0", "classifier_contract": "h0-all-saved-gates-directed-per-leg-frame-v1", "classifier_version": 1, "closure": {"count_matches_source": true, "weight_matches_source": true}, "count_by_status": {"frame_unknown": 1, "near": 2805, "outer": 23254, "proven_out_of_frame": 4278}, "per_leg_scope_rule": "outbound_and_return_each_lte_declared_scope_radius", "source_count": 30338, "source_weight": 775582.0, "weight_by_status": {"frame_unknown": 29.0, "near": 97581.750000000006, "outer": 568530.0833333333, "proven_out_of_frame": 109441.16666666669}, "weight_semantics": "pre_sampling_occupied_housing_candidate_weight"}. Sampled-frame accounting: {"count_by_status": {"artifact_error": 0, "frame_unknown": 0, "in_frame": 1000, "proven_out_of_frame": 0}, "rate_by_status": {"artifact_error": 0.0, "frame_unknown": 0.0, "in_frame": 100.0, "proven_out_of_frame": 0.0}, "source_count": 1000, "source_weight": 97581.749999999853, "weight_by_status": {"artifact_error": 0.0, "frame_unknown": 0.0, "in_frame": 97581.749999999853, "proven_out_of_frame": 0.0}}.

Modeling caveats: 1. corridor routing is a single-reviewer accuracy check (~89.5%; see the report's Appendix O) 2. primary H0 unresolved weight is 3.800% 3. fixed-gate access-policy interval spans 17.300 ppQA and artifact warnings: None [L41]

Twenty-Mile Primary Scope

Publication state: Modeled estimate (3 caveats).

The 20-mile primary H0 channel contains 2,000 draws on the full in-frame sampled denominator: 49 unobserved (2.6%), 1,607 confirmed exposure (75.3%), 243 possible exposure (16.1%), 101 unresolved model (5.9%). The separate robust H0/H1/H2 lattice reports 39 unobserved (2.2%), 1,400 confirmed exposure (67.0%), 248 possible exposure (15.7%), 313 unresolved model (15.1%). Neither channel drops or renormalizes unresolved-model weight.

H0 uncertainty intervals: direction-H0 2.6% to 24.7%; fixed-H0-gate access 2.2% to 17.7%; joint 2.2% to 33.0%; spatial/grade-H0 1.8% to 7.0%. These are distinct sensitivity results and are not averaged.

Candidate-frame accounting: {"all_saved_gates_evaluated": true, "certificate_artifact_path": "frame_classification.json", "certificate_artifact_sha256": "36907288f4cf7a4ccfabd54f0ee709028695812e693a2a1e915dcc25e6f092c0", "classifier_contract": "h0-all-saved-gates-directed-per-leg-frame-v1", "classifier_version": 1, "closure": {"count_matches_source": true, "weight_matches_source": true}, "count_by_status": {"frame_unknown": 1, "near": 2805, "outer": 23254, "proven_out_of_frame": 4278}, "per_leg_scope_rule": "outbound_and_return_each_lte_declared_scope_radius", "source_count": 30338, "source_weight": 775582.0, "weight_by_status": {"frame_unknown": 29.0, "near": 97581.750000000006, "outer": 568530.0833333333, "proven_out_of_frame": 109441.16666666669}, "weight_semantics": "pre_sampling_occupied_housing_candidate_weight"}. Sampled-frame accounting: {"count_by_status": {"artifact_error": 0, "frame_unknown": 0, "in_frame": 2000, "proven_out_of_frame": 0}, "rate_by_status": {"artifact_error": 0.0, "frame_unknown": 0.0, "in_frame": 100.0, "proven_out_of_frame": 0.0}, "source_count": 2000, "source_weight": 666111.8333333333, "weight_by_status": {"artifact_error": 0.0, "frame_unknown": 0.0, "in_frame": 666111.8333333333, "proven_out_of_frame": 0.0}}.

Modeling caveats: 1. corridor routing is a single-reviewer accuracy check (~89.5%; see the report's Appendix O)2. primary H0 unresolved weight is 5.934%3. fixed-gate access-policy interval spans 15.508 pp**QA and artifact warnings:** None [L41]

Relationship Between the Two Scopes

20-mile minus 5-mile weighted unobserved-rate difference: primary conditional H0 **+0.4** percentage points; separate robust H0/H1/H2 lattice **+0.6** percentage points. These channels answer different questions and are not averaged. Each directed outbound and return leg must independently satisfy its declared 5/20-road-mile scope; the total round trip may approach 10/40 miles. An origin with no unchanged saved gate qualifying under H0 is outside that scope before sampling [L41].

References

ID	Source and use in this report
E20	U.S. Supreme Court, Carpenter v. United States opinion, https://www.supremecourt.gov/opinions/17pdf/16-402_h315.pdf - Digital location privacy policy/legal context.

ID	Source and use in this report
E21	Library of Congress, <i>United States v. Jones</i> , 565 U.S. 400, https://www.loc.gov/item/usrep565400/ - GPS tracking/ location surveillance policy/legal context.
E37	Flock Group Inc., U.S. Patent No. 11,416,545 B1, “System and Method for Object-Based Query of Video Content Captured by a Dynamic Surveillance Network,” issued Aug. 16, 2022, https://patents.google.com/patent/US11416545B1/en - Flock’s own patent. Confirmed language: a “dynamic surveillance network” that is device-agnostic and ingests video from “unrelated surveillance devices” (neighborhood, residential, business, traffic, portable/mobile cameras); a “dynamic surveillance and object-based query” (DSOQ) server/system; object-based query of video “based on content of the video footage and not just time and location”; attribute profiling of people (“male, female, race,” height/weight/clothing) and vehicles (make/model/color/plate); and “face recognition data points”/facial-similarity feature-vector matching. Used only for the vendor’s own described architecture and intended capability; the report does not claim Flock has deployed any specific patented function, including facial recognition. Accessed 2026-07-13.
E44	U.S. Government Accountability Office, <i>Military Housing: DOD Should Address Critical Supply Shortages</i> (GAO-25-106208), product page, https://www.gao.gov/products/gao-25-106208 - GAO summary states that about two-thirds of service members in the U.S. live off base in local communities . Used for national off-base housing share context for force-protection commute exposure. Not an installation-level census. Accessed 2026-07-16.
E45	Congressional Research Service, <i>Military Housing</i> (CRS Report R47728), https://www.congress.gov/crs-product/R47728 - CRS reports that about 58% of servicemembers receive a tax-free housing allowance (BAH) and obtain housing from the local housing market outside their installation. Complements E44 (slightly different metric: BAH recipients vs “live off base”). Used for off-base housing / civilian-market commute framing. Accessed 2026-07-16.

ID	Source and use in this report
E46	DoD Travel Management Office / Defense Travel Management Office, Basic Allowance for Housing overview, https://www.travel.dod.mil/Allowances/Basic-allowance-for-Housing/ - Official BAH framing: compensation based on local civilian housing markets when government quarters are not provided; members choose off-base housing in the civilian market. Used only for policy mechanism (BAH → civilian infrastructure), not for a percentage. Accessed 2026-07-16.
E49	U.S. Department of Defense, DoD Manual 4165.63, <i>DoD Housing Management</i> , incorporating Change 3, April 1, 2025, Enclosure 3, section 1.c.(2), https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodm/416563m.pdf - Official housing-market-analysis guidance: market-area boundaries may use a 20-mile or one-hour peak-traffic commute, with military/public information and installation-command recommendations supporting local adjustment. Used to justify a predeclared 20-road-mile primary reachability scope, not to claim that most personnel at a particular installation live within that distance. Accessed 2026-07-20.
E104	U.S. Navy, Naval Base San Diego, "About San Diego," https://cnrsw.cnic.navy.mil/Installations/NAVBASE-San-Diego/Resources/About-San-Diego/ - Official description of Naval Base San Diego as the Navy's premier Pacific Fleet surface-force installation, its homeported ships, tenant commands, Navy Region Southwest headquarters, medical complex, and family housing. Published ship and tenant counts are dated site context, not an operational schedule. Accessed 2026-07-29.
E105	U.S. Navy, Naval Base San Diego, "Mission and Vision," https://cnrsw.cnic.navy.mil/Installations/NAVBASE-San-Diego/About/Mission-and-Vision/ - Current official mission and vision emphasizing installation support, resilience, and operational readiness for the Fleet, Fighters, and Families. Used as public mission framing without inferring ship status, sailing time, or personnel assignments. Accessed 2026-07-29.

ID	Source and use in this report
E148	United Kingdom National Cyber Security Centre, "UK and allies expose Russian intelligence campaign targeting western logistics and technology organisations," May 21, 2025, https://www.ncsc.gov.uk/news/uk-partners-expose-russian-intelligence-campaign - Joint-government attribution that GRU Unit 26165 targeted logistics and transportation organizations and internet-connected cameras near military installations and border crossings to monitor and track aid shipments. Accessed 2026-08-08.
E149	U.S. Government Accountability Office, <i>Information Security: OPM Has Improved Controls, but Further Efforts Are Needed</i> (GAO-17-614), August 3, 2017, https://www.gao.gov/products/gao-17-614 - Official account that the 2015 OPM breaches affected 21.5 million individuals and involved sensitive personnel and security-clearance-investigation information. Accessed 2026-08-08.
E150	Federal Bureau of Investigation, "The Threat Posed by the Chinese Government and the Chinese Communist Party to the Economic and National Security of the United States," July 7, 2020, https://www.fbi.gov/news/speeches-and-testimony/the-threat-posed-by-the-chinese-government-and-the-chinese-communist-party-to-the-economic-and-national-security-of-the-united-states - FBI attribution and counterintelligence context for Chinese theft of OPM and other large U.S. personal-data repositories. Accessed 2026-08-08.
E171	Naval Sea Systems Command, Southwest Regional Maintenance Center, "About," https://www.navsea.navy.mil/Home/RMC/SWRMC/About/ - Official description of SWRMC's contribution to maintaining Navy ships and systems. Accessed 2026-08-08.
E172	U.S. Navy, "Capt. Brian Karosich Assumes Command of Southwest Regional Maintenance Center," July 12, 2024, https://www.navy.mil/Press-Office/News-Stories/Article/3840176/capt-brian-karosich-assumes-command-of-southwest-regional-maintenance-center/ - Official confirmation of SWRMC activity at Naval Base San Diego. Accessed 2026-08-08.

ID	Source and use in this report
L2	Supporting Tools and Docs/ supporting_memoranda/ FlockSafety_Four_Vulnerability_Chain.md ; Supporting Tools and Docs/ supporting_documents/finding_cards/ Finding_01_... - Finding_04_... ; source-evidence: Raw Evidence/Documentation/ My_Four_Vulnerabilities.md - Vulnerability timeline, Cure53-style finding cards, chain framing, remediation status.
L4	Supporting Tools and Docs/ supporting_documents/analyses/ FlockNova_Application_Capability_Enumeration. md - Nova search, access-control, and client-side enforcement-gap findings.
L5	Supporting Tools and Docs/ supporting_documents/analyses/ FlockSafety_Dev_Subdomain_Bundle_Enumeration. md - 370-host headline footprint, dev bundle families, and repeated-key evidence.
L8	Raw Evidence/FlockSafety/ flock_hardcoded_api_check.json - Primary validation artifact for key scope and 50 item grants.
L11	Supporting Tools and Docs/ supporting_documents/analyses/ ArcGIS_Search_App_Data.md - Search app data model, people-search device pool, plate/vehicle/search workflows.
L12	Supporting Tools and Docs/ supporting_documents/analyses/ ArcGIS_Hotlist_App_Data.md - Hotlist, alerts, hit history, and alerting data model.
L13	Supporting Tools and Docs/ supporting_documents/analyses/ ArcGIS_Mapped_Attributes.md - FlockOS/map layers and attributes; CAD/AVL, patrol/mobile, body-camera, drone, Raven/audio, and Flock911 surfaces; Flock911 incident-layer state, audio/transcript token classes, per- word start/end timing, selected-event state, and playback/ scrub controls.

ID	Source and use in this report
L14	Raw Evidence/ArcGIS/flocksafety_prod_arcgis_data/analysis_summary.json; source dump: Raw Evidence/ArcGIS/flocksafety_prod_arcgis_data/features_*.json - 2025-12-14 ArcGIS export summary: 335,701 records, status/device-family counts from the package feature dump.
L15	Raw Evidence/ArcGIS/flocksafety_prod_arcgis_data/count_verification.json - Field verification: total/status breakdown; active=1 (214,233) vs status=inService (214,073) split; confirms scale claims use status, not active.
L20	Raw Evidence/Documentation/My_Four_Vulnerabilities.md; P1 - Repository disclosure record and published article for my 53-exposure research, Finding 1 scope analysis, and public ArcGIS corroboration method.
L35	Raw Evidence/Screenshots/permissions_dronetoo.png; Raw Evidence/FlockNova/nova_flockdemo_com/js/nova_flockdemo_com_App.js - Permission vocabulary including canCaravanSearch, Aerodome access, device tunnels, AVL sharing, nationwide lookup, and ArcGIS sign-in.
L37	Raw Evidence/FlockSafety/dev_app_flocksafety_com/dev-beefeater; Raw Evidence/FlockSafety/dev_app_flocksafety_com/prod-beefeater - Burp XML wire records for Finding 3: anonymous GETs to Beefeater /api/v1/token/esri and /token/google returning HTTP 200 with tokens; no cookie/authorization.

ID	Source and use in this report
L40	Superseded earlier detailed run; the full per-origin paired study records are retained by the researcher offline for authorized review and are not shipped in this package. - The granular per-origin records (synthetic origins, gates, blocks, selected and attempted routes, blocked-camera observations, outcome records, seeds, and summary hashes) come from an earlier method run. They were removed from the shipped package because they were superseded by the V7 run and contain exact per-origin routes near the installations. The canonical V7 route-coverage estimates and their SHA-256 summary bindings are in L41 and control any transcription conflict.
L41	Complete V7 snapshot and staged cohort <code>reachability-v7-audit-92bd2e6a45c3a7c6</code> (index SHA-256 <code>3d0755cb6fcc7dd165c6b9dd53ade8bfe5402773231904ea3807d1f17dc3b96b</code>); primary conditional H0 and robust H0/H1/H2 results, methodology, frame accounting, split uncertainty, QA flags and warnings, hashes, and annexes. Canonical source: the hash-bound V7 snapshot at <code>../Supporting Tools and Docs/project_support/planning_and_qa/reachability_v7_audit_development_snapshot.json</code> .

Per-site records

The V7 result for this installation - primary conditional H0 and robust H0/H1/H2 outcomes, frame accounting, uncertainty intervals, QA flags, seeds, and SHA-256 summary bindings - is recorded in the hash-bound V7 snapshot [L41]. The large per-origin staging artifacts (route geometry, frame classification, review packets) are not shipped in this package; the full records are retained by the researcher offline for authorized review [L40].