

Pentagon

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

The Pentagon is the headquarters of the Department of Defense — the Secretary of Defense, the Joint Staff, and the military departments all work here. A late-night gathering of Joint Staff, communications, continuity, and security traffic could be an early sign the building had gone into crisis response.

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 57.2% of those modeled trips pass a camera. Another 22.3% might pass one, and 20.0% could not be determined. Closer in, at 5 miles, 34.3% of trips could not even be fully routed. This is a model of routes, not proof that anyone was watched.

HEADLINE RESULT · MODELED TRIPS TO A GATE

20 MILES OUT



■ **57.2%**

Passes a camera

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

■ **22.3%**

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°.

■ **0.5%**

Avoids all cameras

No confirmed camera on any plausible road.

■ **20.0%**

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



■ **52.0%** Passes ■ **13.3%** Might pass ■ **0.4%** Avoids ■ **34.3%** Couldn't determine

Couldn't determine is 34.3% here because the router often could not complete a legal round trip to a reviewed gate — typically missing road links, one-way restrictions, or no qualifying return.

Installation Context

The Pentagon is the headquarters of the Department of Defense and the principal workplace for the Office of the Secretary of Defense, the Joint Staff, military-department leadership, defense agencies, and supporting organizations [E77-E79]. The Joint Chiefs of Staff advise the President, Secretary of Defense, and National Security Council, while the Pentagon Reservation also depends on continuous security, continuity, emergency-management, facilities, and communications support [E78, E79, E174, E175].

The installation's intelligence value therefore lies less in ordinary building occupancy than in rare convergence. Senior civilian and military leaders, Joint Staff directorates, combatant-command representatives, interagency officials, foreign-defense delegations, communications personnel, protective staff, and continuity organizations may use the same public approaches even though a single vehicle reveals none of those roles. The National Military Command Center is publicly acknowledged within the Pentagon complex, but public sources do not disclose its watch schedules, crisis procedures, or current activity [E175].

The study uses these public facts only to construct defensive scenarios. It did not identify a Pentagon employee, leader, vehicle, delegation, protective detail, meeting, or crisis. The route model covers four reviewed access points and does not establish which entrance any real person uses.

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 Joint Staff or National Military Command Center Crisis Activation

Assume that a foreign intelligence service obtains a sufficiently long history of vehicle observations on public roads serving the Pentagon and gradually associates recurring plates with probable Joint Staff personnel, senior-defense offices, communications support, security staff, continuity personnel, and liaison organizations. The service would not need a complete roster. Repetition could establish ordinary weekday, nighttime, weekend, and exercise baselines while identifying vehicles that recur during publicly visible crises.

The highest-value signal would be a sustained convergence across normally separate communities. Suppose vehicles probably associated with current-operations staff, senior military advisers, communications and command-support personnel, continuity organizations, and protective or facility-security functions begin arriving outside established patterns and remain through the night. A single

observation would prove little. A multi-day activation involving several correlated clusters could indicate that the Pentagon's national-level command and advisory system was responding to an emerging contingency rather than conducting routine business [E77-E79, E174, E175].

The adversary might not know the underlying decision, military option, or geographic focus. It could nevertheless estimate when the response began, whether activity was escalating, and which external agencies or defense organizations appeared involved. Correlation with diplomatic statements, military movements, cyber events, or overseas developments could provide early warning before the United States publicly acknowledged a change in posture.

The personnel risk is equally serious. Once vehicles are associated with senior leaders, advisers, communications personnel, or protective staff, historical observations may expose probable residences, family vehicles, recurring off-site meetings, and locations where official security is reduced. Movement data can therefore become both an organizational-warning source and a targeting or coercion file.

A direct precedent was reported by the Associated Press in 2026. According to AP sources, Israel exploited Iran's street- and traffic-camera infrastructure to build movement and relationship intelligence on senior officials and their protective personnel, including routes, parking locations, associated addresses, and occasions when multiple officials were present together [E141]. That report is not evidence of compromise at the Pentagon. It demonstrates that a state will repurpose civilian surveillance infrastructure to locate senior leadership and detect rare convergence when the strategic value is high.

The present study does not show that a Pentagon crisis activation occurred or that any leader was identified. It demonstrates Tier A route exposure around a national command site. The inference would require Tier B historical observations, Tier C identity or organizational enrichment, a defensible baseline, and careful treatment of the Pentagon's large unrelated workforce and visitor population.

Scenario 2: Mapping a Sensitive Defense Delegation or Interagency Decision Cell

A second risk arises from visitors rather than permanent personnel. The Pentagon routinely supports senior-defense leadership, military departments, defense agencies, interagency coordination, and foreign-defense engagement [E77-E79, E174]. A hostile service could use repeated vehicle co-occurrence to identify rental vehicles, embassy-associated transportation, government fleets, contractor vehicles, or recurring liaison personnel that appear together during a narrow series of visits.

Suppose several unfamiliar vehicles begin arriving during the same period as a stable cluster associated with one Joint Staff or senior-defense community. Those vehicles are also observed at hotels, airports, embassies, contractor offices, or other government facilities. Over time, the analyst may infer the size and duration of a delegation, which U.S. organizations are meeting it, whether senior leadership joined the engagement, and whether the pattern concerns a one-time visit or a sustained planning cell.

The subject of the meeting could remain unknown. Yet the identity of the participating country, agency, service, or industry partner can itself reveal a sensitive defense relationship, contingency-planning effort, acquisition review, coalition consultation, or response to a developing crisis. The intelligence service could then direct cyber collection, human-source approaches, or surveillance toward the smaller group of likely participants rather than the Pentagon's entire workforce.

The Royal Concierge program provides a close precedent. Reporting based on Snowden documents described a GCHQ system that monitored hotel reservations to identify the timing and location of diplomats' travel and to cue additional technical or human collection [E147]. Its value did not depend on reading the substance of a diplomatic meeting. It came from identifying rare overlaps among officials and knowing when and where the overlap would occur.

A longitudinal ALPR archive could supply a similar warning layer around the Pentagon. Vehicles would not prove attendance, diplomatic status, or meeting purpose, and shared roads, drivers, hotels, ride services, and neighboring destinations would create substantial alternative explanations. But repeated, cross-source association could narrow a population of thousands to a small group warranting further collection.

The study did not obtain visitor logs, plate histories, embassy records, hotel data, or identity enrichment. It therefore cannot identify a delegation or meeting. The scenario explains why routine vehicle observations around a national decision center become more sensitive when retained, searchable, and combinable with other travel and identity systems.

Combined Significance

The first scenario concerns strategic warning: an adversary could detect that national defense leadership and command-support organizations had entered an abnormal operating posture. The second concerns association intelligence: the adversary could identify who is participating in a sensitive consultation before learning what the participants are discussing. Together they show why a plate observation near the Pentagon can acquire intelligence value through repetition even when no message content is collected.

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 (4 caveats)	Modeled estimate (4 caveats)
Headline channel	robust_lattice	robust_lattice

Field	5-mile near-base	20-mile primary
Primary H0 U/C/P/R	1.1%/65.9%/16.8%/16.2%	0.6%/66.3%/24.8%/8.3%
Robust H0/H1/H2 U/C/P/R	0.4%/52.0%/13.3%/34.3%	0.5%/57.2%/22.3%/20.0%
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 **-0.5** percentage points; separate robust H0/H1/H2 lattice **+0.1** percentage points. These channels answer different questions and are not averaged. [L41]

Legislative Significance

The Pentagon illustrates why national-level command sites require controls beyond ordinary law-enforcement privacy rules. Legislators should require heightened vendor security, short non-hit retention, case-specific historical searches, restricted external sharing, immutable audit logs, and prompt notification to Pentagon force-protection and counterintelligence authorities when an ALPR credential, repository, or vendor environment affecting the National Capital Region is compromised. Those controls should be designed without publishing exact camera gaps, preferred approaches, or gate-level avoidance information.

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

Finding	Observed condition	Evidence boundary	Disclosure and remediation status	Source(s)
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
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 counterintelligence and force-protection concern arises because personnel, contractors, and visitors associated with the Pentagon must use public transportation infrastructure outside the protected facility. The route model does not establish where any real employee or visitor lives or how any individual travels. It evaluates only the public-road exposure of the declared synthetic residential frame.

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 Pentagon 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 Pentagon 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 `pentagon`. It binds cohort `reachability-v7-audit-92bd2e6a45c3a7c6`, pair `pentagon-f0c3002779c25f88`, seed `382605111`, index SHA-256 `3d0755cb6fcc7dd165c6b9dd53ade8bfe5402773231904ea3807d1f17dc3b96b`, five-mile summary SHA-256 `8c856d5dc5f90044ead4570fd6f3f73dd1353247992bc25fc3c99ab74526702d`, and twenty-mile summary SHA-256 `88025d463bd743523e4eac23bb59cb384c4a745815fa4cf055d6a001c18dfa81`. [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 (4 caveats)	Modeled estimate (4 caveats)
V7 cohort	<code>reachability-v7-audit-92bd2e6a45c3a7c6</code>	<code>reachability-v7-audit-92bd2e6a45c3a7c6</code>
Pair ID	<code>pentagon-f0c3002779c25f88</code>	<code>pentagon-f0c3002779c25f88</code>
Pair seed	<code>382605111</code>	<code>382605111</code>
Generated at	<code>2026-08-12T15:54:22.709631+00:00</code>	<code>2026-08-12T15:54:22.709631+00:00</code>
Summary SHA-256	<code>8c856d5dc5f90044ead4570fd6f3f73dd1353247992bc25fc3c99ab74526702d</code>	<code>88025d463bd743523e4eac23bb59cb384c4a745815fa4cf055d6a001c18dfa81</code>
Routing engine	<code>valhalla-3.6.3-auto-frozen-osm-camera-blind-directed-overlay-v6</code>	<code>valhalla-3.6.3-auto-frozen-osm-camera-blind-directed-overlay-v6</code>
Synthetic draws	1000	2000

Field	5-mile near-base	20-mile primary
Headline channel	robust_lattice	robust_lattice
Primary H0 U/C/P/R counts	11/659/168/162	16/1323/434/227
Primary H0 U/C/P/R	1.1%/65.9%/16.8%/16.2%	0.6%/66.3%/24.8%/8.3%
Robust H0/H1/H2 U/C/P/R counts	4/520/133/343	9/1104/376/511
Robust H0/H1/H2 U/C/P/R	0.4%/52.0%/13.3%/34.3%	0.5%/57.2%/22.3%/20.0%
Direction H0 interval	1.1% to 34.1%	0.6% to 33.7%
Fixed-gate access interval	0.5% to 35.6%	0.5% to 20.7%
Joint interval	0.4% to 48.0%	0.5% to 42.8%
Spatial H0 interval	1.1% to 1.4%	0.6% to 2.1%
Gate reselection diagnostic	fixed 0.6%; reselected 0.5%; absolute difference 0.1 pp	fixed 0.5%; reselected 0.5%; absolute difference 0.0 pp
H0 gate near-tie sensitivity	60s: 32.1% disagreement weight; 120s: 40.7% disagreement weight; 300s: 57.5% disagreement weight	60s: 47.4% disagreement weight; 120s: 59.6% disagreement weight; 300s: 71.0% disagreement weight
H0 range sensitivity U/C/P/R	75ft: 9.4%/65.9%/8.5%/16.2%; 150ft: 1.1%/65.9%/16.8%/16.2%; 250ft: 0.0%/66.1%/17.7%/16.2%	75ft: 9.8%/66.1%/15.8%/8.3%; 150ft: 0.6%/66.3%/24.8%/8.3%; 250ft: 0.3%/66.4%/25.0%/8.3%
H0 direction sensitivity U/C/P/R	25deg: 1.2%/65.9%/16.7%/16.2%; 30deg: 1.2%/65.9%/16.7%/16.2%; 35deg: 1.1%/65.9%/16.8%/16.2%; 40deg: 1.1%/65.9%/16.8%/16.2%; 45deg: 1.1%/65.9%/16.8%/16.2%; 60deg: 1.1%/65.9%/16.8%/16.2%	25deg: 0.7%/66.3%/24.7%/8.3%; 30deg: 0.6%/66.3%/24.8%/8.3%; 35deg: 0.6%/66.3%/24.8%/8.3%; 40deg: 0.6%/66.3%/24.8%/8.3%; 45deg: 0.6%/66.3%/24.8%/8.3%; 60deg: 0.6%/66.3%/24.8%/8.3%
Deployment-status sensitivity U/C/P/R	declared_policy: 1.1%/65.9%/16.8%/16.2%; in_service_only: 1.8%/65.4%/16.6%/16.2%	declared_policy: 0.6%/66.3%/24.8%/8.3%; in_service_only: 1.1%/65.9%/24.7%/8.3%
Physical-camera-cluster influence	maximum removal shift 8.2 pp; 2 material clusters	maximum removal shift 6.4 pp; 1 material cluster
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 16.200%3. fixed-gate access-policy interval spans 35.100 pp4. one or more physical camera clusters exceed the influence materiality limit	1. corridor routing is a single-reviewer accuracy check (~89.5%; see the report's Appendix O)2. primary H0 unresolved weight is 8.283%3. fixed-gate access-policy interval spans 20.163 pp4. one or more physical camera clusters exceed the influence materiality limit

Field	5-mile near-base	20-mile primary
QA and artifact warnings	None	None

20-mile minus 5-mile weighted unobserved-rate difference: primary conditional H0 **-0.5** percentage points; separate robust H0/H1/H2 lattice **+0.1** percentage points. These channels answer different questions and are not averaged. [L41]

Reachability Results

Five-Mile Near-Base Scope

Publication state: Modeled estimate (4 caveats).

The 5-mile primary H0 channel contains 1,000 draws on the full in-frame sampled denominator: 11 unobserved (1.1%), 659 confirmed exposure (65.9%), 168 possible exposure (16.8%), 162 unresolved model (16.2%). The separate robust H0/H1/H2 lattice reports 4 unobserved (0.4%), 520 confirmed exposure (52.0%), 133 possible exposure (13.3%), 343 unresolved model (34.3%). Neither channel drops or renormalizes unresolved-model weight.

H0 uncertainty intervals: direction-H0 1.1% to 34.1%; fixed-H0-gate access 0.5% to 35.6%; joint 0.4% to 48.0%; spatial/grade-H0 1.1% to 1.4%. 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":

"61cc2bd53787e2a2356ac845b183ae6f5fdb3deb8c6d94da33729d7493331dbd", "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": 483, "near": 7098, "outer": 52482, "proven_out_of_frame": 40382},

"per_leg_scope_rule": "outbound_and_return_each_lte_declared_scope_radius", "source_count":

100445, "source_weight": 1796322.0000000007, "weight_by_status": {"frame_unknown":

6387.333333333334, "near": 195482.75, "outer": 868215.1666666674, "proven_out_of_frame":

726236.7499999999}, "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": 195482.74999999744, "weight_by_status": {"artifact_error": 0.0,

"frame_unknown": 0.0, "in_frame": 195482.74999999744, "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 16.200% 3. fixed-gate access-policy interval spans 35.100 pp 4. one or more physical camera clusters exceed the influence materiality limit **QA and artifact warnings:** None [L41]

Twenty-Mile Primary Scope

Publication state: Modeled estimate (4 caveats).

The 20-mile primary H0 channel contains 2,000 draws on the full in-frame sampled denominator: 16 unobserved (0.6%), 1,323 confirmed exposure (66.3%), 434 possible exposure (24.8%), 227 unresolved model (8.3%). The separate robust H0/H1/H2 lattice reports 9 unobserved (0.5%), 1,104 confirmed exposure (57.2%), 376 possible exposure (22.3%), 511 unresolved model (20.0%). Neither channel drops or renormalizes unresolved-model weight.

H0 uncertainty intervals: direction-H0 0.6% to 33.7%; fixed-H0-gate access 0.5% to 20.7%; joint 0.5% to 42.8%; spatial/grade-H0 0.6% to 2.1%. 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":

"61cc2bd53787e2a2356ac845b183ae6f5fdb3deb8c6d94da33729d7493331dbd", "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": 483, "near": 7098, "outer": 52482, "proven_out_of_frame": 40382},

"per_leg_scope_rule": "outbound_and_return_each_lte_declared_scope_radius", "source_count":

100445, "source_weight": 1796322.0000000007, "weight_by_status": {"frame_unknown":

6387.333333333334, "near": 195482.75, "outer": 868215.1666666674, "proven_out_of_frame":

726236.7499999999}, "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": 1063697.9166666661, "weight_by_status": {"artifact_error": 0.0, "frame_unknown":

0.0, "in_frame": 1063697.9166666661, "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 8.283% 3. fixed-gate access-policy interval spans 20.163 pp 4. one or more physical camera clusters exceed the influence materiality limit **QA and artifact warnings:** None [L41]

Relationship Between the Two Scopes

20-mile minus 5-mile weighted unobserved-rate difference: primary conditional H0 **-0.5** percentage points; separate robust H0/H1/H2 lattice **+0.1** 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, <i>Carpenter v. United States</i> opinion, https://www.supremecourt.gov/opinions/17pdf/16-402_h315.pdf - Digital location privacy policy/legal context.
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.

ID	Source and use in this report
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.
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.
E77	U.S. Department of Defense, “About,” Pentagon Tours, https://www.defense.gov/About/Pentagon-Tours/ - Official statement that the Pentagon is Department of Defense headquarters and that DoD provides military forces to deter war and ensure national security. Used for institutional context, not for current occupancy or schedules. Accessed 2026-07-21.

ID	Source and use in this report
E78	U.S. Department of Defense, Office of the Director of Administration and Management, "About," https://dam.defense.gov/About-DA-M/index.html - Official description of Pentagon Reservation governance, administrative support, continuity, security, law-enforcement, and emergency-management responsibilities supporting senior defense leadership and headquarters tenants. Accessed 2026-07-21.
E79	Washington Headquarters Services, "Facilities Services Directorate," https://www.whs.mil/About-WHS/Directorates/Facilities-Services-Directorate-FSD/ - Official public context for continuous Pentagon Reservation facility readiness and support to Office of the Secretary of Defense, Joint Staff, defense agencies, and military departments across the National Capital Region. Accessed 2026-07-21.
E141	Associated Press via PBS NewsHour, "Iran built a vast camera network to control dissent. Israel used it to track targets, AP sources say," March 23, 2026, https://www.pbs.org/newshour/world/iran-built-a-vast-camera-network-to-control-dissent-israel-used-it-to-track-targets-ap-sources-say - Reported precedent for state exploitation of street and traffic cameras to map senior officials, protective personnel, routes, parking locations, and meetings. Accessed 2026-08-08.
E147	The Guardian, "Snowden cache reveals diplomats' hotel bookings being tracked by GCHQ," November 17, 2013, https://www.theguardian.com/world/2013/nov/17/edward-snowden-diplomats-hotel-bookings-documents - Reported Royal Concierge precedent for monitoring diplomats' hotel reservations to identify travel timing and location and cue additional technical or human intelligence collection. Accessed 2026-08-08.
E174	Joint Chiefs of Staff, "About the Joint Chiefs of Staff," https://www.jcs.mil/about/ - Official description of the Chairman and Joint Chiefs as military advisers to the President, Secretary of Defense, and National Security Council. Accessed 2026-08-08.
E175	U.S. Department of Defense, "National Military Command Center," https://www.defense.gov/Multimedia/Photos/igphoto/2002932553/ - Official public acknowledgment of the National Military Command Center in the Pentagon. Used only to establish the publicly acknowledged facility, not its procedures or current activity. 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].