Jet Propulsion Laboratory

National Security Exposure Assessment

Joshua Michael
Email: josh@nexanet.ai

September 2, 2026

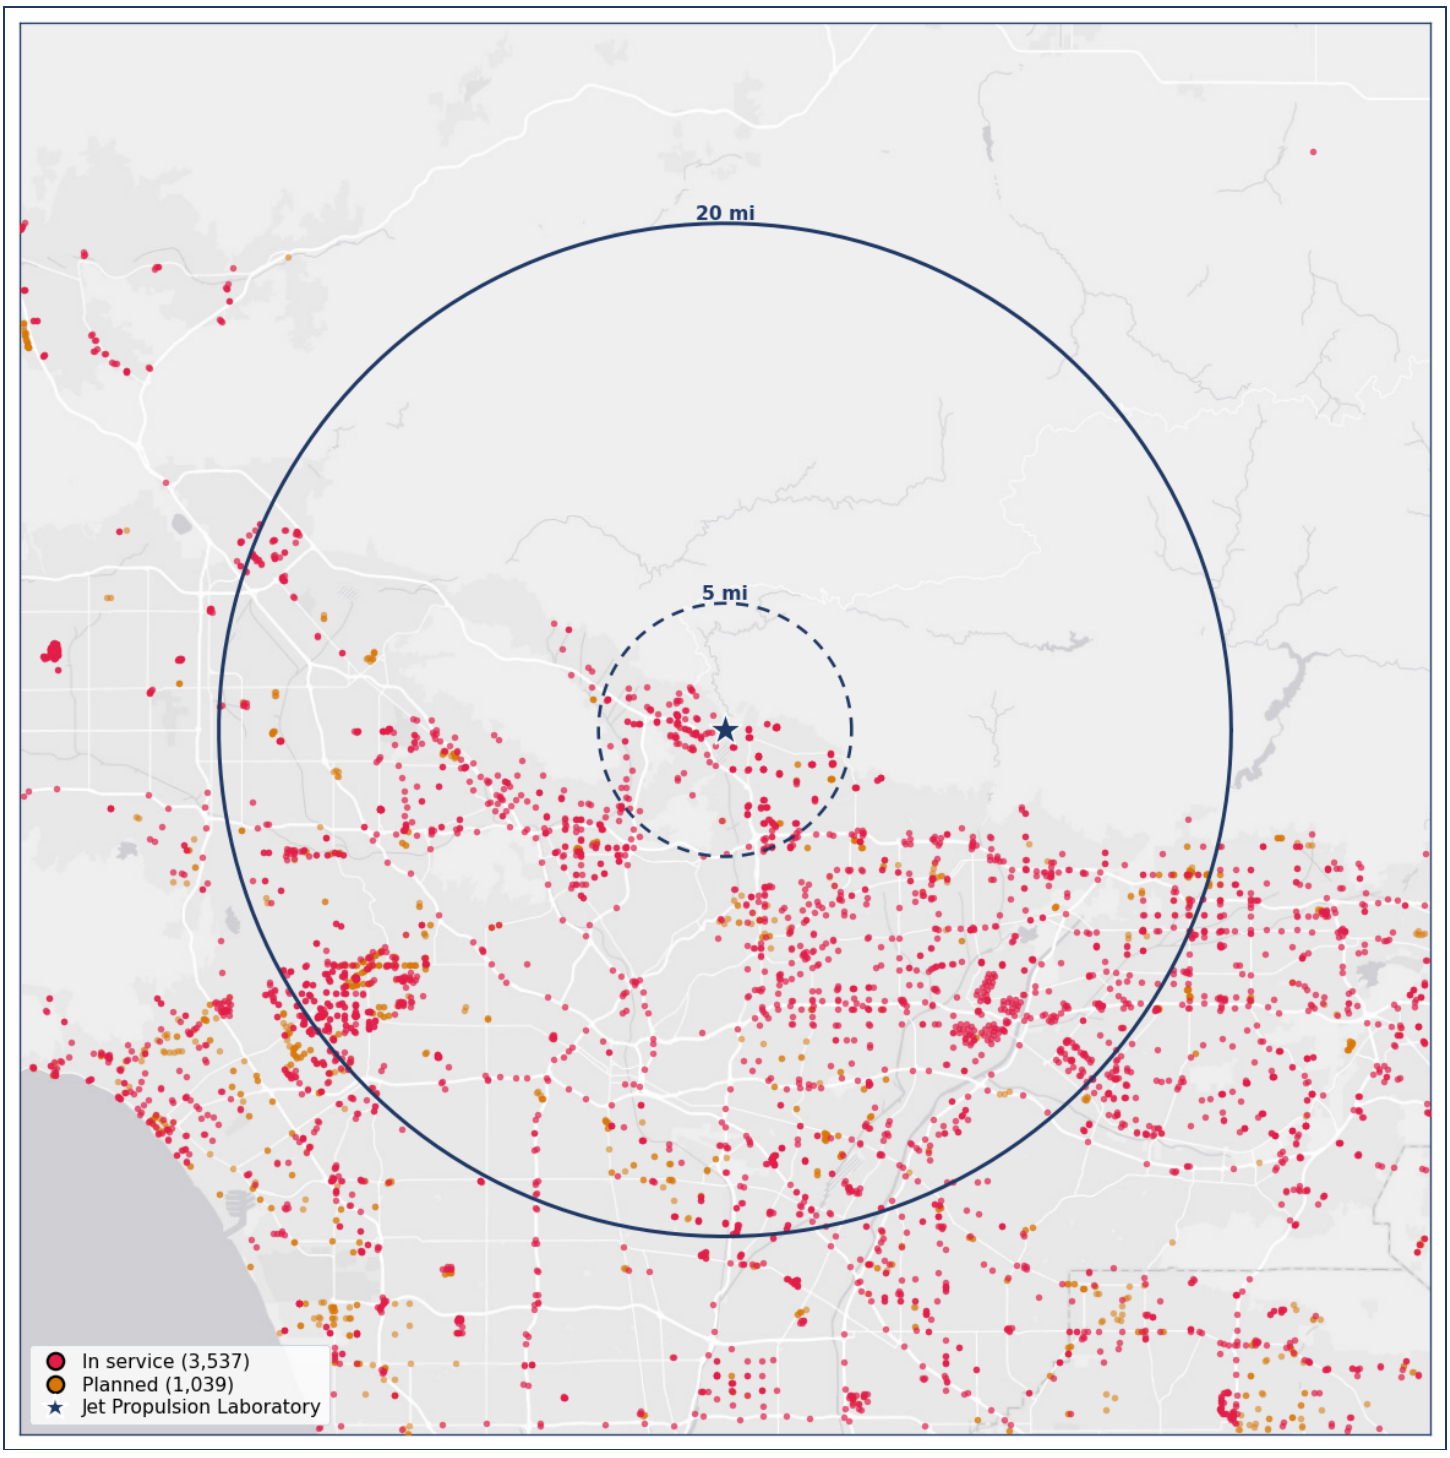

Executive Assessment

JPL in Pasadena is NASA’s robotic-exploration lab. It runs the Deep Space Network that talks to spacecraft across the solar system, and a Rover Operations Center for Moon and Mars missions. A sudden late-night gathering of mission controllers, navigators, and network staff could mean a spacecraft was in trouble — before NASA said so.

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 87.9% of those modeled trips pass a camera. Another 3.9% might pass one, 1.6% avoid them, and 6.7% could not be determined. Even at 5 miles, about 85.5% 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

87.9% — **Passes a camera**
A confirmed camera sits on every plausible road to a gate.

3.9% — **Might pass**
A camera is nearby, but facing could not be confirmed. The file's rotationAngle is a housing-orientation field, not travel direction, so unnamed or poorly oriented cameras are not counted as hits.

1.6% — **Avoids all cameras**
No confirmed camera on any plausible road.

6.7% — **Couldn't determine**
The model could not finish a legal round trip for these draws.

5 MILES OUT

85.5% Passes **1.6%** Might pass **9.7%** Avoids **3.2%** Couldn't determine

Installation Context

The Jet Propulsion Laboratory is a federally funded research and development center managed by Caltech for NASA and is a leader in robotic space exploration [E120]. JPL manages NASA's Deep Space Network, the international system that tracks, commands, and receives data from deep-space spacecraft, and operates mission-control and ground-systems functions supporting many missions [E121]. In 2025, JPL also opened a Rover Operations Center intended to support current and future Moon, Mars, aerial, and surface missions and external partnerships [E122].

These functions depend on specialized, often small technical communities: mission controllers, navigation teams, flight-software engineers, spacecraft-system experts, communications and network personnel, robotics specialists, security staff, and partner organizations. Their movement can become informative when it deviates from an established baseline, even though a single vehicle observation says nothing about a mission.

The study did not identify a JPL employee, mission, spacecraft, partner, anomaly, vehicle, or household. JPL personnel use multiple work arrangements and external facilities, and the route model cannot determine who occupies a vehicle or why it approaches the laboratory.

Laboratory-Specific Mission-Assurance Scenarios

The scenarios require Tier B historical observations and Tier C identity or organizational enrichment beyond the Tier A deployment and route information used in the study.

Scenario 1: Detecting a Spacecraft Anomaly or Critical Mission Event Through After-Hours Team Activation

Assume that a foreign intelligence service obtains longitudinal vehicle observations around JPL and identifies recurring clusters associated with mission controllers, navigation specialists, flight-software engineers, Deep Space Network personnel, project leadership, and security or communications staff. The analyst would establish the normal rhythm for public mission milestones and ordinary operations, then search for deviations that do not match the public schedule.

A sudden nighttime or weekend convergence could be highly informative. If a small technical group that normally works separately begins arriving together, remains for extended periods, and is joined by senior mission leadership or outside partner personnel, the pattern could indicate a spacecraft anomaly, critical trajectory maneuver, communications loss, recovery effort, or other mission event. The adversary might not know which spacecraft is affected, but public DSN schedules, mission status pages, launch calendars, and partner travel could narrow the possibilities.

Early knowledge of a mission problem could support cyber targeting against the affected project, disinformation, attempts to interfere with recovery communications, or collection against partner organizations. Repetition could also reveal which specialists are called during particular categories of failure, creating a human map of mission dependencies.

JPL has a direct same-site compromise precedent. NASA's Office of Inspector General reported that, in 2011, intruders using Chinese-based IP addresses gained full access to eighteen servers supporting key JPL missions, including the Deep Space Network, and could copy, modify, or delete sensitive files and mission-critical accounts; 87 gigabytes of data were uploaded to the attackers' addresses [E123]. The report did not establish in public that the intrusion was directed by the Chinese government, so this assessment does not make that attribution. It does establish that foreign-origin intruders gained deep access to the same mission environment whose external staffing patterns are at issue here.

The route study did not observe a mission activation or identify any JPL specialist. The inference would require a long baseline, reliable role attribution, and independent mission corroboration. Shift work, scheduled events, facility maintenance, and unrelated projects remain alternative explanations.

Scenario 2: Identifying Deep Space Network and Rover-Operations Specialists and Their Partner Network

A second scenario focuses on the technical population rather than an event. Assume that an analyst identifies vehicles that repeatedly overlap with Deep Space Network operations, the Rover Operations Center, mission-ground-system groups, and visiting university, industry, or foreign-space-agency partners. Historical co-occurrence could reveal which individuals bridge projects, which outside organizations recur, and which specialists are called to several missions.

The result would be a prioritized counterintelligence list. A person who repeatedly appears during DSN emergencies, critical maneuvers, or rover-operations events may possess knowledge of network limitations, scheduling conflicts, mission anomalies, ground-system architecture, or recovery procedures. Once linked to a probable identity and residence, that person can be targeted through phishing, credential theft, recruitment, travel surveillance, or device compromise away from JPL's protected systems.

Partner mapping creates a second layer of risk. Repeated co-arrival by an outside organization and a particular JPL technical group can reveal a relationship before it is publicly described. Even when the project is civil, the same people and technologies may support export-controlled, dual-use, or national-security-relevant work.

The suspected Chinese intelligence compromise of Marriott's Starwood reservation system provides a direct association-data precedent. The intrusion persisted for years and exposed travel and identity records; investigators cited by Reuters said the operation appeared designed for intelligence collection

rather than financial gain [E124]. Hotel and reservation data were valuable because they could reveal who traveled, where, and when. A longitudinal vehicle archive can perform the same relationship-mapping function at a mission laboratory.

The JPL study did not identify a DSN or rover-operations specialist, an external partner, or a real commute. The scenario requires historical records, accurate identity and employer enrichment, and corroboration. Visitors, carpools, rentals, shared roads, and unrelated projects create significant uncertainty.

Combined Significance

The first scenario concerns mission warning: roadway activity could reveal that a spacecraft or communications problem has triggered an emergency response. The second concerns technical-personnel and partner identification. Together, they show how a commercial movement archive can expose both the moment a space mission is under stress and the small group of people best positioned to explain or resolve it.

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 (2 caveats)	Modeled estimate (3 caveats)
Headline channel	robust_lattice	robust_lattice
Primary H0 U/C/P/R	10.2%/87.7%/0.9%/1.2%	1.7%/91.7%/3.9%/2.7%
Robust H0/H1/H2 U/C/P/R	9.7%/85.5%/1.6%/3.2%	1.6%/87.9%/3.9%/6.7%
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 **-8.5** percentage points; separate robust H0/H1/H2 lattice **-8.1** percentage points. These channels answer different questions and are not averaged. [L41]

Legislative Significance

JPL demonstrates that sensitive-site ALPR policy must include mission-critical civilian and federally funded research facilities, not only military bases. The laboratory's role in spacecraft control, the Deep Space Network, robotics, and international mission partnerships means that longitudinal movement records can create mission-assurance and counterintelligence risks even though the campus is operated by a civilian institution.

Legislators should consider heightened security, retention, sharing, audit, and breach-notification obligations for ALPR coverage around federally funded research and development centers and other critical space infrastructure. Notification pathways should include NASA and appropriate federal counterintelligence authorities, while public reporting should avoid disclosing exact route gaps or technical staff patterns.

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 counterintelligence and mission-assurance concern arises because employees, contractors, government counterparts, visitors, and partners use public roads outside the site's controlled perimeter. The occupied-housing frame is a standardized residential-distribution proxy; it does not identify where anyone associated with Jet Propulsion Laboratory lives, establish an employment relationship, or estimate actual commute volume.

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 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

To preserve comparability across the multi-site assessment, the analysis applies the same five-road-mile near-site scope and twenty-road-mile primary scope used for military installations. The twenty-mile distance was originally selected from DoD Manual 4165.63 as a transparent commuting-area

approximation [E49]. At Jet Propulsion Laboratory, it functions only as a standardized sensitive-site catchment. It is not a claim about employee residences, commute norms, company or laboratory policy, or the location of any real worker.

Five miles isolates the immediate approach area. Twenty miles adds the broader occupied-housing stratum while retaining every near-site draw and route unchanged. Reporting both shows how the modeled observation surface changes with a wider road network without converting the sample into an employee or visitor estimate.

Limitations and Non-Claims

This is a model over dated deployment, Census, and road snapshots. Occupied housing is a residential-distribution proxy, not a Jet Propulsion Laboratory workforce or visitor roster; it does not weight origins by employer, role, clearance, shift, measured traffic, public transit use, remote work, or actual residential location.

Camera location and status fields can be incomplete, stale, planned, duplicated, moved, or differently interpreted. OpenStreetMap roads and address points can omit valid travel or include obsolete topology. Gate or access-point availability and site geometry depend on analyst review. The route engine models corridor exposure and a geometric field of view; it does not model optical character recognition probability, weather, plate condition, device uptime, 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. Shared roads, neighboring destinations, rental vehicles, ride services, carpools, public transit, and unrelated visitors create substantial alternative explanations at a corporate or research headquarters.

Evidence and Reproducibility

The V7 drafting packet was built from the complete validated cohort for `jet_propulsion_laboratory`. It binds cohort `reachability-v7-audit-92bd2e6a45c3a7c6`, pair `jet_propulsion_laboratory-ca7d2d45d9c49dbd`, seed `1254176254`, index SHA-256 `3d0755cb6fcc7dd165c6b9dd53ade8bfe5402773231904ea3807d1f17dc3b96b`, five-mile summary SHA-256 `596a01ce33c7e692ac81a2afad5e6987e7aa96a96d651a4f745df1e9eb62346f`, and twenty-mile summary SHA-256 `8ff6fdf35df7c82215b00687ce82e62e0d41b1a69f6da22a16d802e69ad5461e`. [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 (2 caveats)	Modeled estimate (3 caveats)
V7 cohort	reachability-v7-audit-92bd2e6a45c3a7c6	reachability-v7-audit-92bd2e6a45c3a7c6
Pair ID	jet_propulsion_laboratory-ca7d2d45d9c49dbd	jet_propulsion_laboratory-ca7d2d45d9c49dbd
Pair seed	1254176254	1254176254
Generated at	2026-08-12T11:29:13.545551+00:00	2026-08-12T11:29:13.545551+00:00
Summary SHA-256	596a01ce33c7e692ac81a2afad5e6987e7aa96a96d651a4f745df1e9eb62346f	8ff6fdf35df7c82215b00687ce82e62e0d41b1a69f6da22a16d802e69ad5461e
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	102/877/9/12	115/1796/49/40
Primary H0 U/C/P/R	10.2%/87.7%/0.9%/1.2%	1.7%/91.7%/3.9%/2.7%
Robust H0/H1/H2 U/C/P/R counts	97/855/16/32	109/1735/56/100
Robust H0/H1/H2 U/C/P/R	9.7%/85.5%/1.6%/3.2%	1.6%/87.9%/3.9%/6.7%
Direction H0 interval	10.2% to 12.3%	1.7% to 8.3%
Fixed-gate access interval	9.7% to 13.5%	1.6% to 8.2%
Joint interval	9.7% to 14.5%	1.6% to 12.1%

Field	5-mile near-base	20-mile primary
Spatial H0 interval	3.8% to 11.3%	0.2% to 4.0%
Gate reselection diagnostic	fixed 9.9%; reselected 9.9%; absolute difference 0.0 pp	fixed 1.7%; reselected 1.7%; absolute difference 0.0 pp
H0 gate near-tie sensitivity	60s: 0.0% disagreement weight; 120s: 0.0% disagreement weight; 300s: 0.0% disagreement weight	60s: 0.0% disagreement weight; 120s: 0.0% disagreement weight; 300s: 0.0% disagreement weight
H0 range sensitivity U/C/P/R	75ft: 10.3%/86.0%/2.5%/1.2%; 150ft: 10.2%/87.7%/0.9%/1.2%; 250ft: 10.2%/87.7%/0.9%/1.2%	75ft: 1.7%/91.7%/3.9%/2.7%; 150ft: 1.7%/91.7%/3.9%/2.7%; 250ft: 1.6%/91.8%/3.9%/2.7%
H0 direction sensitivity U/C/P/R	25deg: 11.1%/87.7%/0.0%/1.2%; 30deg: 10.5%/87.7%/0.6%/1.2%; 35deg: 10.2%/87.7%/0.9%/1.2%; 40deg: 10.2%/87.7%/0.9%/1.2%; 45deg: 10.2%/87.7%/0.9%/1.2%; 60deg: 10.2%/87.7%/0.9%/1.2%	25deg: 1.7%/91.6%/3.9%/2.7%; 30deg: 1.7%/91.7%/3.9%/2.7%; 35deg: 1.7%/91.7%/3.9%/2.7%; 40deg: 1.7%/91.7%/3.9%/2.7%; 45deg: 1.7%/91.7%/3.9%/2.7%; 60deg: 1.7%/91.7%/3.9%/2.7%
Deployment-status sensitivity U/C/P/R	declared_policy: 10.2%/87.7%/0.9%/1.2%; in_service_only: 10.2%/87.7%/0.9%/1.2%	declared_policy: 1.7%/91.7%/3.9%/2.7%; in_service_only: 1.7%/91.5%/4.1%/2.7%
Physical-camera-cluster influence	maximum removal shift 1.8 pp; 0 material clusters	maximum removal shift 0.9 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 1.200%	1. corridor routing is a single-reviewer accuracy check (~89.5%; see the report's Appendix O)2. primary H0 unresolved weight is 2.733%3. fixed-gate access-policy interval spans 6.675 pp
QA and artifact warnings	None	None

20-mile minus 5-mile weighted unobserved-rate difference: primary conditional H0 **-8.5** percentage points; separate robust H0/H1/H2 lattice **-8.1** percentage points. These channels answer different questions and are not averaged. [L41]

Reachability Results

Five-Mile Near-Base Scope

Publication state: Modeled estimate (2 caveats).

The 5-mile primary H0 channel contains 1,000 draws on the full in-frame sampled denominator: 102 unobserved (10.2%), 877 confirmed exposure (87.7%), 9 possible exposure (0.9%), 12 unresolved model (1.2%). The separate robust H0/H1/H2 lattice reports 97 unobserved (9.7%), 855 confirmed exposure (85.5%), 16 possible exposure (1.6%), 32 unresolved model (3.2%). Neither channel drops or renormalizes unresolved-model weight.

H0 uncertainty intervals: direction-H0 10.2% to 12.3%; fixed-H0-gate access 9.7% to 13.5%; joint 9.7% to 14.5%; spatial/grade-H0 3.8% to 11.3%. 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": "51513b3ae30fe9e04fdd475da63dcff94d94261a10ad74f61d08401a4dc78ba7", "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": 0, "near": 3974, "outer": 59803, "proven_out_of_frame": 112870}, "per_leg_scope_rule": "outbound_and_return_each_lte_declared_scope_radius", "source_count": 176647, "source_weight": 2409397.0000000014, "weight_by_status": {"frame_unknown": 0.0, "near": 40012.74999999998, "outer": 921914.749999999, "proven_out_of_frame": 1447469.5000000026}, "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": 40012.75000000095, "weight_by_status": {"artifact_error": 0.0, "frame_unknown": 0.0, "in_frame": 40012.75000000095, "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 1.200%**QA 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: 115 unobserved (1.7%), 1,796 confirmed exposure (91.7%), 49 possible exposure (3.9%), 40 unresolved model (2.7%). The separate robust H0/H1/H2 lattice reports 109 unobserved (1.6%), 1,735 confirmed exposure (87.9%), 56 possible exposure (3.9%), 100 unresolved model (6.7%). Neither channel drops or renormalizes unresolved-model weight.

H0 uncertainty intervals: direction-H0 1.7% to 8.3%; fixed-H0-gate access 1.6% to 8.2%; joint 1.6% to 12.1%; spatial/grade-H0 0.2% to 4.0%. 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": "51513b3ae30fe9e04fdd475da63dcff94d94261a10ad74f61d08401a4dc78ba7", "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": 0, "near": 3974, "outer": 59803, "proven_out_of_frame": 112870}, "per_leg_scope_rule": "outbound_and_return_each_lte_declared_scope_radius", "source_count": 176647, "source_weight": 2409397.0000000014, "weight_by_status": {"frame_unknown": 0.0, "near": 40012.74999999998, "outer": 921914.749999999, "proven_out_of_frame": 1447469.5000000026}, "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": 961927.4999999974, "weight_by_status": {"artifact_error": 0.0, "frame_unknown": 0.0, "in_frame": 961927.4999999974, "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 2.733%3. fixed-gate access-policy interval spans 6.675 pp**QA and artifact warnings:** None [L41]

Relationship Between the Two Scopes

20-mile minus 5-mile weighted unobserved-rate difference: primary conditional H0 **-8.5** percentage points; separate robust H0/H1/H2 lattice **-8.1** 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, 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, United States v. Jones, 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.
E49	U.S. Department of Defense, DoD Manual 4165.63, *DoD Housing Management*, 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.
E120	NASA Jet Propulsion Laboratory, “Who We Are,” https://www.jpl.nasa.gov/who-we-are/ - Official description of JPL as a NASA-funded, Caltech-managed leader in robotic space exploration and its specialized workforce. Accessed 2026-08-08.
E121	NASA Jet Propulsion Laboratory, “Deep Space Network,” https://www.jpl.nasa.gov/missions/dsn/ and https://www.jpl.nasa.gov/news/nasas-deep-space-network-turns-60-and-prepares-for-the-future/ - Official DSN mission and JPL management of spacecraft tracking, commanding, and data return. Accessed 2026-08-08.

ID	Source and use in this report
E122	NASA Jet Propulsion Laboratory, “NASA JPL Unveils Rover Operations Center for Moon, Mars Missions,” December 10, 2025, https://www.jpl.nasa.gov/news/nasa-jpl-unveils-rover-operations-center-for-moon-mars-missions/ - Official mission and partnership context for the Rover Operations Center. Accessed 2026-08-08.
E123	NASA Office of Inspector General, “Cybersecurity Management and Oversight at the Jet Propulsion Laboratory,” Report IG-19-022, June 18, 2019, https://oig.nasa.gov/docs/IG-19-022.pdf - Official account of the 2011 intrusion involving Chinese-based IP addresses and full access to servers supporting key JPL missions, including the DSN; no public state attribution is asserted here. Accessed 2026-08-08.
E124	Reuters, “Exclusive: Clues in Marriott hack implicate China - sources,” December 5, 2018, https://www.reuters.com/article/world/exclusive-clues-in-marriott-hack-implicate-china-sources-idUSKBN1O504R/ - Reporting that investigators believed the long-running Starwood reservation compromise was designed for Chinese intelligence collection; used as precedent for travel and association mapping. Accessed 2026-08-08.
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.

ID	Source and use in this report
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.
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.

ID	Source and use in this report
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.
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 `reachability-v7-audit-92bd2e6a45c3a7c6` (index SHA-256 `3d0755cb6fcc7dd165c6b9dd53ade8bfe5402773231904ea3807d1f17dc3b96b`); 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 `../Supporting Tools and Docs/project_support/planning_and_qa/reachability_v7_audit_development_snapshot.json`.

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].