RTX Corporation Headquarters

National Security Exposure Assessment

Joshua Michael
Email: josh@nexanet.ai

© 2026 Joshua Michael · CC BY 4.0. Use with credit.

September 2, 2026

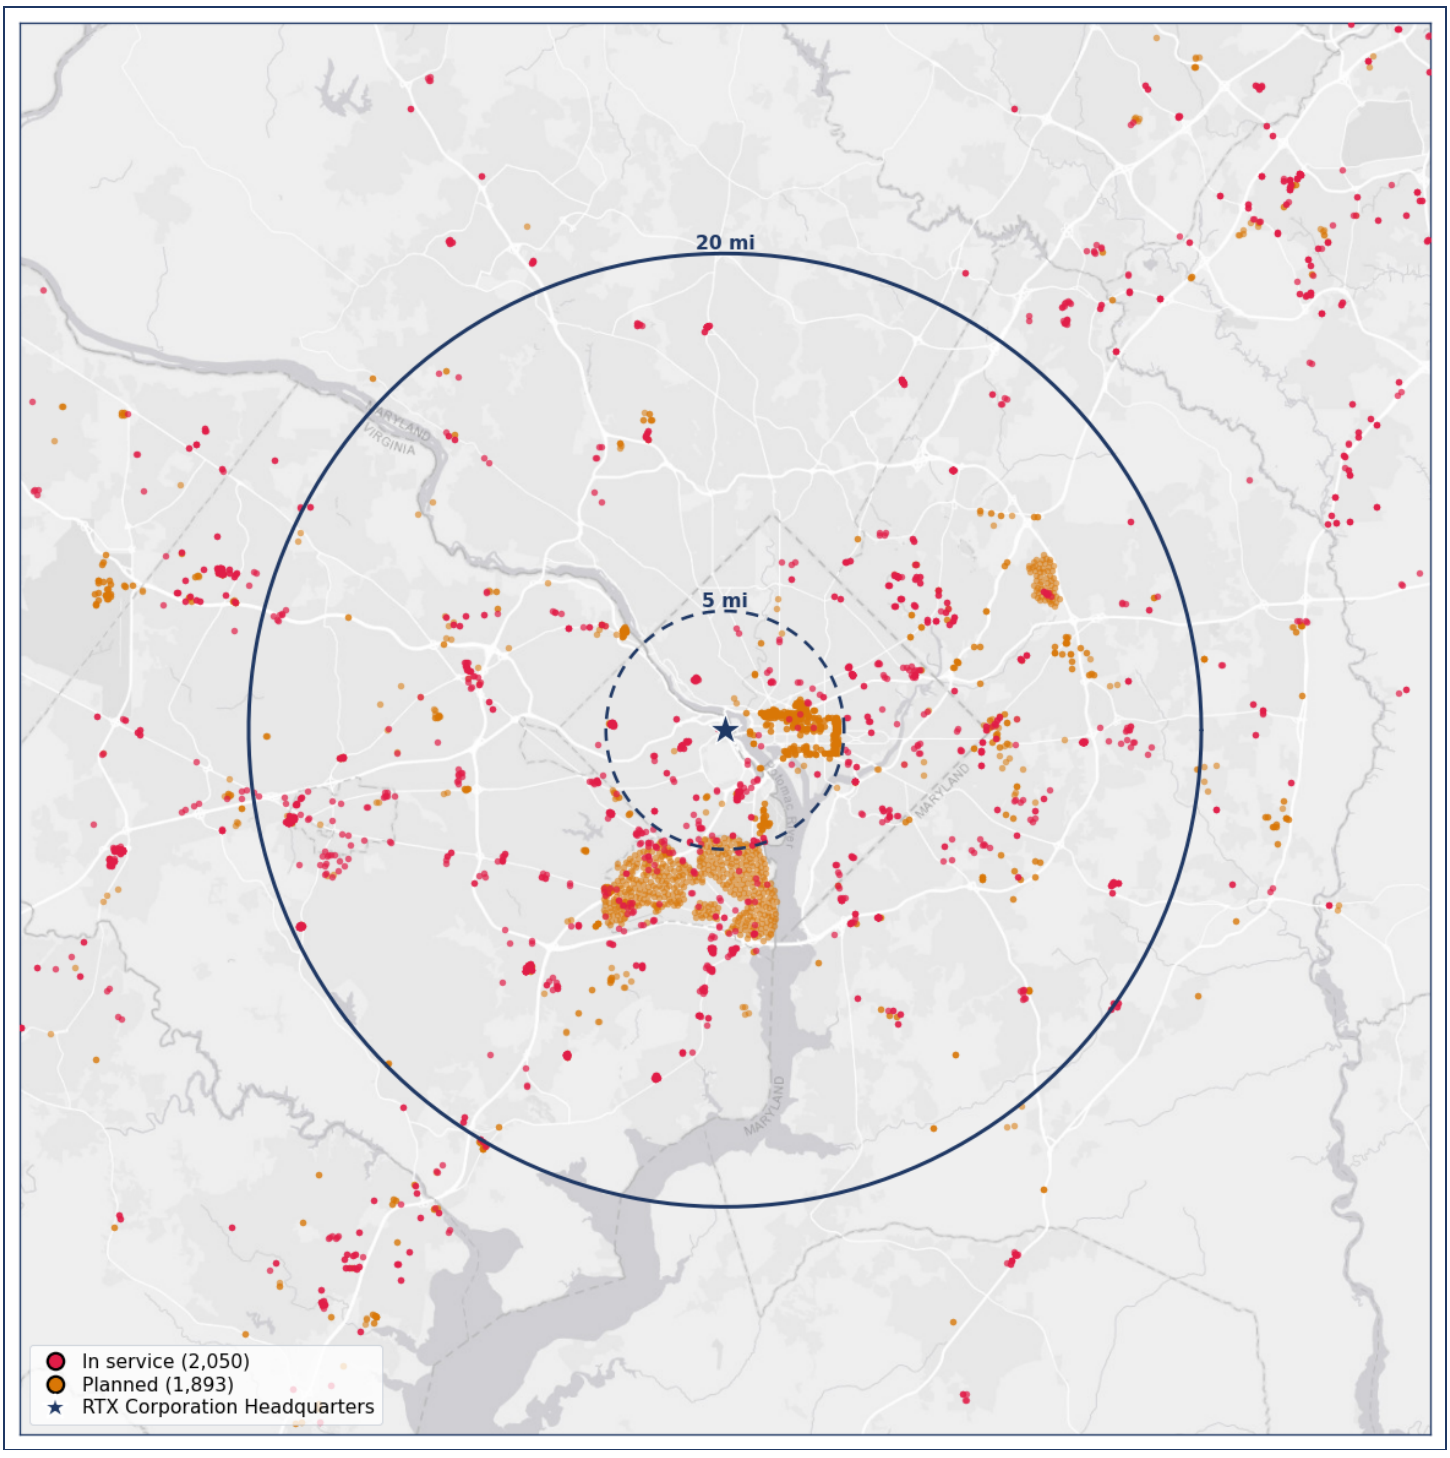

Executive Assessment

RTX’s headquarters in Arlington is the corporate office for Raytheon, Collins, and Pratt & Whitney. Patriot and NASAMS air-defense systems used by the U.S. and allies are reviewed from here even though the factories are elsewhere. U.S. officials, allied visitors, and Raytheon program people showing up together could mark a Patriot or NASAMS review — and which partner nation is in the room.

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 55.9% of those modeled trips pass a camera. Another 27.4% might pass one, 4.0% avoid them, and 12.7% could not be determined. Closer in, at 5 miles, 22.4% 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

55.9% **Passes a camera**
A confirmed camera sits on every plausible road to a gate.

27.4% **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°.

4.0% **Avoids all cameras**
No confirmed camera on any plausible road.

12.7% **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

70.4% Passes **2.8%** Might pass **4.4%** Avoids **22.4%** Couldn't determine

Couldn't determine is 22.4% here because extra driveway or side-road options could not be fully routed, so those trips stay unresolved instead of being guessed.

Installation Context

RTX is headquartered in Arlington, Virginia and operates through Collins Aerospace, Pratt & Whitney, and Raytheon [E111]. The corporate headquarters is not a missile-production plant, radar test range, or engine factory. The assessment therefore does not infer that ordinary headquarters traffic corresponds directly to production activity at distributed operating locations.

The headquarters is nevertheless sensitive because RTX's Raytheon business publicly supports integrated air and missile defense systems used by the United States and allied governments. Patriot is fielded by nineteen nations, while NASAMS is used by the United States and numerous international partners [E112, E113]. Rare co-location of corporate leaders, government officials, partner-nation representatives, suppliers, legal or export-control personnel, and program executives can reveal a relationship or crisis that is not apparent from any one participant's travel.

The study did not identify an RTX employee, government official, allied delegation, supplier, meeting, or program. Dense National Capital Region traffic, shared roads, public transit, ride services, rental vehicles, and neighboring destinations create substantial confounders.

Headquarters-Specific Counterintelligence Scenarios

The scenarios require historical vehicle observations and reliable identity, employer, or delegation enrichment beyond the Tier A device and route information used in the study.

Scenario 1: Identifying a U.S.–Allied Patriot or NASAMS Program Delegation

Assume that a foreign intelligence service monitors historical vehicle observations around RTX headquarters and can reliably associate selected vehicles with Raytheon leadership, U.S. air-defense officials, embassy or allied-government personnel, outside counsel, security staff, and executives arriving from program locations. The service would search for rare overlaps rather than routine employee traffic.

Suppose vehicles associated with one partner government repeatedly converge with a small Raytheon and U.S. government group. If the pattern appears around a public procurement announcement, combat loss, software or interceptor upgrade, export decision, or production negotiation, the analyst could infer that a Patriot or NASAMS issue involving that country had reached senior review. The movement record would not reveal the negotiating position or technical problem, but it could identify the government involved, the approximate timing, and the officials or corporate leaders most likely to possess the details.

That association map could cue diplomatic collection, cyber targeting, surveillance of off-site meetings, or recruitment approaches against a supplier or adviser with weaker security. Repetition could reveal which individuals bridge a foreign government, U.S. acquisition officials, and RTX leadership.

Royal Concierge is a direct precedent. Documents disclosed by Edward Snowden described a GCHQ program that monitored hotel bookings to identify the timing and location of foreign diplomats' travel and prepare further technical or human collection [E114]. Its value came from detecting rare official presence and overlap, not from the reservation content itself.

The Arlington model did not identify a delegation or program review. The scenario requires accurate enrichment and corroboration, and it remains subject to major alternative explanations such as shared roads, rental vehicles, embassy transportation, and unrelated meetings.

Scenario 2: Detecting an Air-Defense Production or Replenishment Crisis

A second scenario focuses on response rather than delegation. Assume that an analyst establishes baseline movement for RTX corporate leadership, Raytheon air-defense executives, supply-chain personnel, government logistics and acquisition officials, and a small set of critical suppliers. A sudden multi-day convergence following a major conflict, high interceptor expenditure, production disruption, quality issue, or urgent allied request could indicate that an air-defense replenishment or allocation problem had escalated to corporate headquarters.

The adversary would not learn stockpile numbers or production output from vehicle records alone. It could nevertheless infer that a shortage, production bottleneck, export-control decision, or priority allocation was under review. That warning could guide cyber collection against suppliers, attempts to disrupt transportation or components, influence operations aimed at allied governments, or monitoring of facilities likely to receive surge investment.

The scenario becomes stronger when movement is correlated with public contract actions, congressional funding, allied announcements, and observable logistics. The headquarters signal can identify that a problem has reached decision-makers; collection against the distributed production network can then seek the technical detail.

The documented GRU Unit 26165 campaign provides a direct precedent for this intelligence objective. The United States, United Kingdom, and partners reported that the GRU targeted western logistics and technology organizations involved in assistance to Ukraine and internet-connected cameras near military installations and border crossings to track aid movement [E110]. The campaign treated supply-chain organizations and physical movement as one intelligence problem.

The RTX study does not show a production crisis or connect any vehicle to Patriot, NASAMS, a supplier, or a government agency. It demonstrates the modeled surveillance surface around the headquarters. Any crisis inference would require longitudinal history, enrichment, and corroboration.

Combined Significance

The first scenario exposes allied relationships and sensitive program negotiations. The second exposes the timing of a production or replenishment problem affecting air and missile defense. Together, they show how a headquarters movement archive can reveal both who is involved in a defense decision and when operational demand has created a vulnerability in the supporting industrial base.

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	7.3%/88.2%/2.6%/1.9%	4.7%/64.5%/28.9%/1.9%
Robust H0/H1/H2 U/C/P/R	4.4%/70.4%/2.8%/22.4%	4.0%/55.9%/27.4%/12.7%
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 **-2.6** percentage points; separate robust H0/H1/H2 lattice **-0.4** percentage points. These channels answer different questions and are not averaged. [L41]

Legislative Significance

RTX Headquarters shows why defense-industrial movement data can carry alliance and readiness consequences. Longitudinal observations can expose which foreign government is meeting with a U.S. air-defense supplier or when a production and replenishment issue has escalated to enterprise leadership.

Legislators should consider sensitive-site controls for major defense headquarters and their surrounding ALPR networks, including minimized non-hit retention, restricted bulk and reverse searches, narrow cross-agency sharing, mandatory auditability, vendor-access limitations, and prompt notification to counterintelligence and industrial-base authorities when credentials or historical datasets are compromised.

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
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 RTX Corporation Headquarters 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 RTX Corporation Headquarters, 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 RTX Corporation Headquarters 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 `rtx_corporation_headquarters`. It binds cohort `reachability-v7-audit-92bd2e6a45c3a7c6`, pair `rtx_corporation_headquarters-629f892af9fab227`, seed `1541629511`, index SHA-256 `3d0755cb6fcc7dd165c6b9dd53ade8bfe5402773231904ea3807d1f17dc3b96b`, five-mile summary SHA-256 `6146b8b5557805808db419d9b1bafc5ccb67bedaa12ad86f21eca99b0a9739e5`, and twenty-mile summary SHA-256 `c632b4d37585eac46270a7f82ecbe039d6995da4cbc021b2560857af28324aad`. [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	rtx_corporation_headquarters-629f892af9fab227	rtx_corporation_headquarters-629f892af9fab227

Field	5-mile near-base	20-mile primary
Pair seed	1541629511	1541629511
Generated at	2026-08-12T07:42:45.424047+00:00	2026-08-12T07:42:45.424047+00:00
Summary SHA-256	6146b8b5557805808db419d9b1bafc5ccb67bedaa12ad86f21eca99b0a9739e5	c632b4d37585eac46270a7f82ecbe039d6995da4cbc021b2560857af28324aad
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	73/882/26/19	115/1476/371/38
Primary H0 U/C/P/R	7.3%/88.2%/2.6%/1.9%	4.7%/64.5%/28.9%/1.9%
Robust H0/H1/H2 U/C/P/R counts	44/704/28/224	83/1232/355/330
Robust H0/H1/H2 U/C/P/R	4.4%/70.4%/2.8%/22.4%	4.0%/55.9%/27.4%/12.7%
Direction H0 interval	7.3% to 11.8%	4.7% to 35.5%
Fixed-gate access interval	4.4% to 27.6%	4.0% to 16.9%
Joint interval	4.4% to 29.6%	4.0% to 44.1%
Spatial H0 interval	2.4% to 7.7%	2.8% to 6.0%
Gate reselection diagnostic	fixed 6.7%; reselected 6.7%; absolute difference 0.0 pp	fixed 4.6%; reselected 4.6%; 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: 7.6%/88.0%/2.5%/1.9%; 150ft: 7.3%/88.2%/2.6%/1.9%; 250ft: 4.2%/88.2%/5.7%/1.9%	75ft: 8.6%/64.1%/25.4%/1.9%; 150ft: 4.7%/64.5%/28.9%/1.9%; 250ft: 0.7%/64.5%/32.9%/1.9%
H0 direction sensitivity U/C/P/R	25deg: 8.6%/88.2%/1.3%/1.9%; 30deg: 7.3%/88.2%/2.6%/1.9%; 35deg: 7.3%/88.2%/2.6%/1.9%; 40deg: 7.3%/88.2%/2.6%/1.9%; 45deg: 7.3%/88.2%/2.6%/1.9%; 60deg: 6.4%/88.2%/3.5%/1.9%	25deg: 5.4%/64.5%/28.2%/1.9%; 30deg: 4.7%/64.5%/28.9%/1.9%; 35deg: 4.7%/64.5%/28.9%/1.9%; 40deg: 4.7%/64.5%/29.0%/1.9%; 45deg: 4.7%/64.5%/29.0%/1.9%; 60deg: 3.7%/64.5%/29.9%/1.9%
Deployment-status sensitivity U/C/P/R	declared_policy: 7.3%/88.2%/2.6%/1.9%; in_service_only: 9.0%/88.2%/0.9%/1.9%	declared_policy: 4.7%/64.5%/28.9%/1.9%; in_service_only: 28.5%/62.1%/7.5%/1.9%
Physical-camera-cluster influence	maximum removal shift 1.9 pp; 0 material clusters	maximum removal shift 1.7 pp; 0 material clusters

Field	5-mile near-base	20-mile primary
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 1.900%3. fixed-gate access-policy interval spans 23.200 pp	1. corridor routing is a single-reviewer accuracy check (~89.5%; see the report's Appendix O)2. primary H0 unresolved weight is 1.900%3. fixed-gate access-policy interval spans 12.903 pp
QA and artifact warnings	None	None

20-mile minus 5-mile weighted unobserved-rate difference: primary conditional H0 **-2.6** percentage points; separate robust H0/H1/H2 lattice **-0.4** 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: 73 unobserved (7.3%), 882 confirmed exposure (88.2%), 26 possible exposure (2.6%), 19 unresolved model (1.9%). The separate robust H0/H1/H2 lattice reports 44 unobserved (4.4%), 704 confirmed exposure (70.4%), 28 possible exposure (2.8%), 224 unresolved model (22.4%). Neither channel drops or renormalizes unresolved-model weight.

H0 uncertainty intervals: direction-H0 7.3% to 11.8%; fixed-H0-gate access 4.4% to 27.6%; joint 4.4% to 29.6%; spatial/grade-H0 2.4% to 7.7%. 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": "e06693713e4e6fb43d9bb98413734ff67d4d3ada6b92ccbcaea144cb19a63dc7", "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": 474, "near": 7959, "outer": 50658, "proven_out_of_frame": 41525}, "per_leg_scope_rule": "outbound_and_return_each_lte_declared_scope_radius", "source_count": 100616, "source_weight": 1804721.0000000014, "weight_by_status": {"frame_unknown": 6047.583333333333, "near": 184788.50000000003, "outer": 863715.1666666674, "proven_out_of_frame": 750169.7500000006}, "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": 184788.49999999718, "weight_by_status": {"artifact_error": 0.0, "frame_unknown": 0.0, "in_frame": 184788.49999999718, "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 1.900%3. fixed-gate access-policy interval spans 23.200 pp**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 (4.7%), 1,476 confirmed exposure (64.5%), 371 possible exposure (28.9%), 38 unresolved model (1.9%). The separate robust H0/H1/H2 lattice reports 83 unobserved (4.0%), 1,232 confirmed exposure (55.9%), 355 possible exposure (27.4%), 330 unresolved model (12.7%). Neither channel drops or renormalizes unresolved-model weight.

H0 uncertainty intervals: direction-H0 4.7% to 35.5%; fixed-H0-gate access 4.0% to 16.9%; joint 4.0% to 44.1%; spatial/grade-H0 2.8% to 6.0%. 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": "e06693713e4e6fb43d9bb98413734ff67d4d3ada6b92ccbcaea144cb19a63dc7", "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": 474, "near": 7959, "outer": 50658, "proven_out_of_frame": 41525}, "per_leg_scope_rule": "outbound_and_return_each_lte_declared_scope_radius", "source_count": 100616, "source_weight": 1804721.0000000014, "weight_by_status": {"frame_unknown": 6047.583333333333, "near": 184788.50000000003, "outer": 863715.1666666674, "proven_out_of_frame": 750169.7500000006}, "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": 1048503.6666666586, "weight_by_status": {"artifact_error": 0.0, "frame_unknown": 0.0, "in_frame": 1048503.6666666586, "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 1.900%3. fixed-gate access-policy interval spans 12.903 pp**QA and artifact warnings:** None [L41]

Relationship Between the Two Scopes

20-mile minus 5-mile weighted unobserved-rate difference: primary conditional H0 **-2.6** percentage points; separate robust H0/H1/H2 lattice **-0.4** 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.
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.

ID	Source and use in this report
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.
E111	RTX, “Who We Are” and 2025 results, https://www.rtx.com/ and https://www.rtx.com/news/news-center/2026/01/27/rtx-reports-2025-results-and-announces-2026-outlook- - Official Arlington headquarters context and public business structure across Collins Aerospace, Pratt & Whitney, and Raytheon. Accessed 2026-08-08.
E112	RTX Raytheon, “Global Patriot,” https://www.rtx.com/raytheon/what-we-do/integrated-air-and-missile-defense/global-patriot-solutions - Official public description of Patriot and its multinational user base. Accessed 2026-08-08.
E113	RTX Raytheon, “NASAMS,” https://www.rtx.com/raytheon/what-we-do/integrated-air-and-missile-defense/nasams - Official public description of NASAMS and U.S./allied operators. Accessed 2026-08-08.
E114	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 - Reporting on Royal Concierge, used as precedent for identifying diplomatic and official travel overlaps. Accessed 2026-08-08.
E110	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 linking espionage against logistics and technology organizations with camera targeting along military routes. 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 `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].