VOL. I · EST. MMXXVI

NORTH PACIFIC
STRATEGY INITIATIVE

TECHNICAL BRIEFING NO. 01 · V1.0 · JUNE 2026 · FOR DISCUSSION

The Verified Sky

Sensing, certainty, and the law of automated airspace awareness — the state of the art, what it costs, how fast it acts, and what a Canadian operator may lawfully do with it.

AUTHOR

Jesse James

ISSUED

11 June 2026 · Victoria, British Columbia

SERIES

Technical Briefing — the engineering substrate beneath the policy architecture

NPSI.CA/TB1

§ 0 ABSTRACT

A middle power that intends to know what is in its sky — over a port, a pipeline corridor, a flight park, or a farm — now has the tools to know it with engineering certainty, at consumer prices, in under a second. This briefing surveys the 2025–2026 state of the art in computer vision and sensing for monitoring a defined volume of airspace; describes the *verification architecture* that converts a probabilistic detection into a near-certain identification by requiring every condition in a configurable ledger to pass; reports realistic detection-to-action latency; maps the automated responses that are lawful for a private Canadian operator against those that are criminal; and offers a deliberately conservative capability forecast at one, two-and-a-half, and five years.

The central findings are three. First, the binding constraint is no longer compute or cost — a capable monitoring node is buildable for under `CAD $1,500` — but the detection of *small, distant targets*, where published accuracy has plateaued. Second, certainty is an architecture, not a model: the systems that work do not trust a classifier, they corroborate it. Third, in Canada the line between lawful and criminal response is bright, statutory, and indifferent to good intentions.

$249
67 INT8-TOPS EDGE BOARD — INFERENCE IS SOLVED AT THE EDGE

35.92%
BEST TRACKER ON REALISTIC TINY THERMAL TARGETS — THE HONEST NUMBER

<1–3 s
FIRST FRAME TO VERIFIED IDENTIFICATION TO DELIVERED ALERT

The briefing is organized in six sections: the sensor problem (§ I); the verification architecture (§ II); latency (§ III); response and the law (§ IV); a conservative forecast to 2031 (§ V); and a staged reference architecture with explicit gates (§ VI). Sources and notes follow at § VII. Vendor claims are identified as vendor claims throughout; benchmark figures are taken from the peer-reviewed record.

§ I THE SENSOR PROBLEM

No single sensor can watch a sky. Each modality answers one question well and goes blind on the others, and the discipline of airspace awareness consists largely of knowing which blindness is being purchased at which price. The table below summarizes the field as deployed in 2025–2026; ranges are indicative, since vendor figures are rarely stated against a defined target and condition set, and should be read as such.[1]

TABLE 1 · SENSING MODALITIES FOR A DEFINED AIRSPACE VOLUME, 0–3,000 M

MODALITY	ANSWERS	INDICATIVE RANGE	BLIND TO	COST CLASS
RGB camera	What it is; evidence-grade imagery	~1–3 km (zoom optics)	Night, fog, glare, sky-saturated backlight; narrow field of view	Consumer–prosumer
Thermal (LWIR)	Presence in darkness and smoke	100s m – several km	Thermal crossover at dawn and dusk; obscured targets	Prosumer–industrial
Radar (micro-Doppler)	Range, bearing, elevation in all weather; rotor-vs-bird	~1–5 km small UAS	Identity, make, operator	Industrial–defence
RF detection	Control link; locates drone and operator; reads Remote ID	to ~5 km	Autonomous, RF-silent, fibre-guided; everything that does not emit	Prosumer–defence
Acoustic array	Propeller and motor signature beyond line of sight	~300–500 m	Wind above ~5 m/s; urban noise; quiet and gliding aircraft	Consumer–industrial
Event camera	Microsecond-latency motion, extreme dynamic range	research-grade	Static objects; not yet turnkey	Prosumer R&D

The empirical reference point for what cheap, distributed sensing achieves at national scale is Ukraine's Zvook acoustic network: detections reportedly reach the armed forces' Delta situational-awareness system within roughly twelve seconds, at a false-positive rate of about 1.6 per cent, from sensors costing on the order of US$500 each.[2] The lesson generalizes: coverage and corroboration, not exquisite individual sensors, are what produce reliable awareness.

The pairing logic follows from the blindness table. Radar and RF do the *searching*; an electro-optical/infrared camera does the *confirming* and produces the evidence. Commercial counter-UAS platforms — Dedrone (acquired by Axon), DroneShield, Fortem — are at bottom fusion engines that correlate detections across modalities and decline to escalate on a single sensor's word.[3] The standing rule across the industry is two-sensor corroboration before escalation.

One consequence deserves emphasis because it inverts the intuition that more sensors are always better. For targets that are large, slow, mostly within visual range, and electromagnetically silent — a hang glider or paraglider over a coastal British Columbia flight park is the canonical case — RF detection contributes nothing, acoustics contribute little, and a camera-first architecture (RGB plus thermal, with radar optional for range and weather) is both sufficient and the correct allocation of budget. The modality must be matched to the target, not to the catalogue.

> *Coverage and corroboration, not exquisite individual sensors, are what produce reliable awareness. The modality must be matched to the target, not to the catalogue.*

§ II CERTAINTY IS AN ARCHITECTURE

A detector is a probability machine. It reports that a region of pixels resembles a class of object with some confidence, and on a good benchmark day that confidence is well calibrated; on a bad day a gull at altitude is a drone and a kite is an aircraft. The systems that work in production do not ask the detector to be certain. They make certainty *structurally*, by requiring an identification to pass every condition in a configurable ledger before any action fires. The pattern is general; the instantiation overleaf is the flight-park case.

Every component of this ledger is available open-source, and the stack is stable enough to name. Detection: the YOLO family remains the deployable workhorse, with transformer detectors competitive at higher compute. On VisDrone — the standard aerial benchmark, whose objects average roughly 36 pixels — deployable models score in the high-30s to ~51 per cent mAP50 depending on architecture and weight class, and sliced inference materially improves small-object recall.[4] Tracking: ByteTrack (80.3 MOTA on MOT17) and BoT-SORT (80.5 MOTA, with camera-motion compensation and a stronger motion model) convert per-frame detections into persistent tracks with tunable persistence buffers.[5] Open-vocabulary detection — alerting on a natural-language description without retraining — is real but young: Grounding DINO leads accuracy, YOLO-World leads speed, and sustaining useful frame rates on edge hardware remains an optimization exercise rather than a default.[6] The rules layer above all of this is ordinary application code.

VERIFICATION LEDGER · TRACK № 0147		ALL CONDITIONS MUST PASS
[✓]	**Class.** Detector confidence for the target class above a calibrated threshold	DETECTOR
[✓]	**Size.** Estimated physical dimension within bounds, from bounding box, optics, and range	GEOMETRY
[✓]	**Altitude.** Height estimate inside the declared envelope	RADAR / GEOMETRY
[✓]	**Kinematics.** Speed and trajectory consistent with the target's flight profile	TRACKER
[✓]	**Window.** Inside declared operating hours and daylight rules	CLOCK
[✓]	**Geofence.** Position inside the defined volume, point-in-polygon	GEOFENCE
[✓]	**Persistence.** Conditions held across N consecutive frames — the single most effective false-positive suppressor	TRACKER
GATE: 7 / 7 PASSED		→ ACTION AUTHORIZED

It is worth stating plainly what the benchmarks say about the hard case, because this is where reputations in this field go to die. On Anti-UAV410, the standard thermal-infrared drone-tracking benchmark, state accuracy has sat in the low-to-mid 60s for three years — 67.7 per cent in 2023, with 2025 methods clustered between 64 and 67.[7] When the field built a benchmark that looks like the real world — CST Anti-UAV: 220 thermal sequences, more than 240,000 annotations, tiny targets in complex scenes, 20 state-of-the-art trackers evaluated — the best method achieved 35.92 per cent state accuracy, against 67.69 per cent on Anti-UAV410.[8]

67.69%
BEST TRACKER, STANDARD BENCHMARK (ANTI-UAV410)

35.92%
BEST TRACKER, REALISTIC TINY TARGETS IN CLUTTER (CST ANTI-UAV)

≈ ½
REAL-WORLD HARD CASE AS A FRACTION OF HEADLINE PERFORMANCE

Performance on realistic small-and-distant targets is roughly half of headline performance. Any system design that does not assume losing and re-acquiring the track is designed for the benchmark, not the sky.

§ III THE SPEED OF NOTICING

Latency in these systems decomposes into three regimes that are routinely conflated in marketing material: model inference, verification, and action.

ms	0.3–1 s	< 1–3 s
INFERENCE — SINGLE-DIGIT TO TENS OF MILLISECONDS PER FRAME ON EDGE SILICON	VERIFICATION — N-FRAME PERSISTENCE AT 30 FPS; THE DELIBERATE COST OF CERTAINTY	SOFT ACTION — FIRST FRAME TO DELIVERED ALERT, SLEW COMMAND, OR API CALL

Inference. A Jetson Orin Nano Super — NVIDIA's US$249 developer board, launched December 2024 at 67 INT8 TOPS — runs current YOLO models at 30–60+ frames per second under TensorRT; dedicated accelerators in the Hailo class run common detectors in under ten milliseconds.[9] Inference is, for practical purposes, solved at the edge.

Verification. The persistence requirement in the ledger is a deliberate latency purchase: holding all conditions across, say, ten to thirty frames costs a third of a second to a second and buys an order-of-magnitude reduction in false alarms. This is the correct trade for almost every civil application.

Action. A soft response — email, SMS, push notification, webhook, an API call into another system — adds network time and lands the full chain at under a second to a few seconds. A pan-tilt-zoom slew-to-cue, in which radar or RF hands a predicted position to a camera that mechanically slews and locks an auto-tracker, is dominated by motor time: hundreds of milliseconds to roughly two seconds on commodity hardware. Physical response — launching an interceptor aircraft — runs to seconds and minutes and belongs, as § IV explains, to a legal category most readers must not enter. Cloud-routed pipelines add hundreds of milliseconds to seconds and are appropriate for notification and archival, not for the verification loop, which belongs at the edge.

> *The persistence requirement is a deliberate latency purchase: a third of a second buys an order-of-magnitude reduction in false alarms.*

§ IV RESPONSE AND THE LAW

What may a private Canadian operator lawfully *do* when the ledger closes? The answer divides cleanly, and the dividing line is statutory.

Lawful and commercially routine

Notification in every form; camera handoff and autonomous optical tracking; spotlight or audible deterrent, provided no aircraft is endangered; timestamped evidence packages of clip, track, and metadata; and ingestion of Remote ID broadcasts as a corroborating ledger condition for cooperative aircraft. On this last point the regulatory ground is moving in the operator's favour: on 8 June 2026 Transport Canada published Notice of Proposed Amendment 2026-005 (CARAC NPA 06-2026: *RPAS — Remote Identification, Community-Based Organizations, and Designated Airspace*), proposing mandatory Remote ID for most drone operations on a performance basis — Broadcast or Network, both on the ASTM F3411 standard — with the comment window open to 9 September 2026.[10] In the United States, Remote ID has been mandatory and enforced since 16 March 2024. If the Canadian proposal proceeds, cooperative-aircraft identification becomes a broadcast to be received rather than an inference to be made, and the strongest single condition a ledger can carry.

Criminal for a civilian, regardless of intent

Radio-frequency jamming and spoofing are prohibited by the Radiocommunication Act: subsection 4(4) bars the installation, use, possession, manufacture, import, or sale of jammers, and paragraph 9(1)(b) bars interference with radiocommunication, with exemptions under section 14 confined to federal bodies such as the RCMP and the Department of National Defence and to narrow authorized pilots.[11] Taking over a drone's control software engages the unauthorized-computer-use provisions of the Criminal Code, sections 342.1 and 342.2. Physically downing a drone is interference with an aircraft under the Aeronautics Act and CARs Part IX.

The interceptor systems that exist — Fortem's DroneHunter, which according to the company began first customer deliveries of its 5.0 version in January 2026 and was selected by the Pentagon's counter-UAS task force under the Replicator-2 initiative; Anduril's Anvil — are defence and government systems, full stop.[12] A Canadian civil operator's response ceiling is detection, tracking, evidence, and notification. That ceiling is set by Parliament, not by engineering.

Privacy is a design input, not an afterthought

Fixed cameras operated in the course of commercial activity engage PIPEDA and, in British Columbia, Alberta, and Québec, the provincial private-sector statutes. The obligations are settled: a reasonable purpose, visible signage, collection limited to what is necessary, a field of view that does not sweep neighbouring properties or public sidewalks, secured storage, scheduled retention and destruction, and a written policy.[13] Audio is the trap: recording private conversations without consent engages the Criminal Code's interception provisions, and the correct engineering answer is to disable microphones at the hardware level. A system pointed at the sky carries low privacy exposure by construction — one more argument, where targets permit, for the camera-first architecture aimed upward.

STANDING CAVEAT

This briefing is a technical and regulatory survey, not legal advice. The counter-drone and Remote ID landscape is in active rulemaking on both sides of the border; an operator deploying any response capability should verify the current state of the law and obtain counsel.

> *A Canadian civil operator's response ceiling is detection, tracking, evidence, and notification. That ceiling is set by Parliament, not by engineering.*

§ V A CONSERVATIVE FORECAST

Forecasts in this field fail in a characteristic way: they extrapolate the hardware curve, which is real, onto the accuracy curve, which is not. The measured baseline for the hardware curve is the Jetson Orin Nano line, which moved from 40 INT8 TOPS at US$499 in March 2023 to 67 INT8 TOPS at US$249 in December 2024 — roughly a three-to-four-fold improvement in compute per dollar in under two years — while headline TOPS figures across the industry are increasingly inflated by precision-format changes (INT8 to INT4 to FP4) rather than by efficiency gains at constant precision.[14] The measured baseline for the accuracy curve is the benchmark plateau documented in § II. Holding both honestly:

TABLE 2 · CONSERVATIVE CAPABILITY FORECAST — WHAT A BUILDER MAY PLAN ON

HORIZON	NEAR-CERTAIN	LIKELY	GENUINELY UNCERTAIN — DO NOT PLAN ON IT
+1 yr mid-2027	Edge compute 15–25% cheaper per unit; a robust RGB+thermal node in the low four figures; YOLO + ByteTrack/BoT-SORT remains the production stack	Open-vocabulary detection usable at modest frame rates on edge boards for coarse alerting	Any material jump in small-and-distant-target accuracy
+2.5 yr end-2028	Continued hardware cost decline; thermal cores cheaper as 12 µm pitch standardizes	On-device vision-language models practical at the edge — "describe the thing to alert on" without retraining; fusion software increasingly commoditized	Tiny-target tracking in clutter; design for re-acquisition, not for an unbroken track
+5 yr 2031	Multi-camera fusion nodes cheap and unremarkable; edge AI market substantially larger on any analyst definition	Natural-language-configurable monitoring as the normal interface; event-camera sensors entering fusion stacks for fast and backlit targets	Whether small-object detection at distance escapes its plateau at all. Treat a breakout as upside, never as the plan

The asymmetry is the finding. Everything around the detector — compute, sensors, fusion, interface — is on a cost curve that favours the builder. The detector's performance on the hardest realistic case is not. Architecture must therefore carry what the model cannot: corroboration across sensors, persistence across frames, graceful loss and re-acquisition of track, and a verification ledger that converts an imperfect detector into a near-certain system.

§ VI A REFERENCE ARCHITECTURE IN FOUR STAGES

For an operator proceeding from nothing to a deployed system, the staging below orders the spend by what each stage retires in risk. Each stage carries an explicit gate; the gate, not the calendar, authorizes the next stage.

TABLE 3 · STAGED BUILD — GATES BEFORE SPEND

STAGE	BUILD	GATE TO ADVANCE
0 · Prove < CAD $1,500	Single RGB camera; Jetson Orin Nano Super or Raspberry Pi 5 with a Hailo accelerator; YOLO + ByteTrack; the full verification ledger in application code; action = email and webhook	Reliable detection of the declared target at the required range, with fewer than one false alert per day
1 · Harden	Add a 640×512 uncooled thermal camera for night, glare, and contrast failure; a PTZ for slew-to-track; BoT-SORT for motion robustness; evidence logging; Remote ID ingestion as a ledger condition	Two-sensor corroboration driving false positives to approximately zero across a full diurnal cycle
2 · Extend	Small commercial radar for beyond-visual range and all-weather cueing; rigorous time synchronization and a single coordinate convention (WGS84 with a local ENU frame) adopted before, not after, fusion; optionally an open-vocabulary model for reconfigurable alerting	Detection range and weather robustness meeting the site's declared requirement
3 · Respond	Soft responses only: alerting, autonomous optical tracking, deterrence that endangers no aircraft, evidence packages, operator dashboard. No jamming, no spoofing, no takeover, no kinetic capability — see § IV	Counsel's review of the response set against the current state of Canadian law

Two engineering notes from the fusion literature merit standing inclusion because they are cheap to honour and expensive to retrofit. First, clock discipline: a two-hundred-millisecond drift between a radar plot and a camera frame is enough to spawn duplicate tracks and defeat corroboration; synchronize from the first commit. Second, coordinate discipline: one geodetic convention, declared once, carried everywhere. Most fusion failures in the field are bookkeeping failures.

§ VII SOURCES AND NOTES

01 Vendor detection-range claims are generally stated without a defined target, aspect, or condition set; international standardization of counter-UAS detection metrics is in progress at the IEC. All ranges in Table 1 should be de-rated accordingly.

02 Figures for the Zvook acoustic detection network — approximately twelve seconds to appearance in the Delta system, a false-positive rate of about 1.6 per cent, and per-sensor cost on the order of US$500 — are as reported by the project's Air Force liaison in Western press coverage, 2023–2025, and are presented as reported rather than independently audited.

03 On fusion architecture and two-sensor corroboration as the industry false-positive control: vendor technical documentation, Dedrone (Axon) DedroneTracker.AI; DroneShield DroneSentry; Fortem SkyDome; and the counter-UAS trade literature, 2024–2026.

04 VisDrone benchmark results, 2024–2025 literature: deployable single-stage models in the high-30s mAP50 (e.g., improved YOLOv8n variants) to approximately 51 per cent mAP50 for heavier 2025 detection-transformer architectures; sliced-inference (SAHI) gains documented in Akyon et al., 2022, and successors.

05 Zhang et al., ByteTrack: Multi-Object Tracking by Associating Every Detection Box, ECCV 2022 (80.3 MOTA, 77.3 IDF1, MOT17); Aharon et al., BoT-SORT: Robust Associations Multi-Pedestrian Tracking, 2022 (80.5 MOTA, 80.2 IDF1, MOT17).

06 Grounding DINO 1.5 Edge (IDEA Research, 2024): 36.2 zero-shot AP on COCO with an edge-optimized backbone; Cheng et al., YOLO-World: Real-Time Open-Vocabulary Object Detection, CVPR 2024: 35.4 AP on LVIS at 52 fps on server-class hardware. Edge deployment above ~10 fps remains an active optimization problem in the 2025 literature.

07 Anti-UAV410 thermal tracking benchmark: Huang et al., IEEE TPAMI 2023 (SiamDT, 67.7% state accuracy), with 2025 methods reporting 64.0–67.0% — a three-year plateau in the band.

08 Xie, B., Zhang, C., Wang, F., Liu, P., Lu, F., Chen, Z., & Hu, W. (2025). CST Anti-UAV: A Thermal Infrared Benchmark for Tiny UAV Tracking in Complex Scenes. ICCV 2025 Workshops; arXiv:2507.23473. 220 sequences, >240k annotations, 20 trackers evaluated; best method 35.92% state accuracy versus 67.69% on Anti-UAV410.

09 NVIDIA Jetson Orin Nano Super Developer Kit, announced 17 December 2024: US$249, 67 INT8 TOPS, 102 GB/s memory bandwidth (NVIDIA developer documentation). Hailo-8 sub-10 ms detector latency per vendor documentation and the Frigate NVR community's published benchmarks.

10 Transport Canada, Notice of Proposed Amendment 2026-005 (CARAC NPA 06-2026): RPAS — Remote Identification, Community-Based Organizations, and Designated Airspace, published 8 June 2026; comment period to 9 September 2026; performance-based Remote ID accepting Broadcast or Network paths on ASTM F3411. United States: 14 CFR Part 89, Remote ID enforcement from 16 March 2024.

11 Radiocommunication Act, R.S.C. 1985, c. R-2, ss. 4(4), 9(1)(b), 14; Criminal Code, R.S.C. 1985, c. C-46, ss. 342.1, 342.2; Aeronautics Act, R.S.C. 1985, c. A-2; Canadian Aviation Regulations, Part IX, as amended by SOR/2025-70 (BVLOS and other operations framework, in force 4 November 2025).

12 Fortem Technologies press release, Lindon, Utah, January 2026 (DroneHunter 5.0 first customer deliveries; selection under the Replicator-2 initiative) — company statements, reported as such. Anduril Industries, Anvil product documentation.

13 Office of the Privacy Commissioner of Canada, Guidelines for Overt Video Surveillance in the Private Sector; PIPEDA Report of Findings #2010-008; Personal Information Protection Act (British Columbia), S.B.C. 2003, c. 63.

SERIES INDEX

WP-01	The Bilateral Foundation — A thirty-year Canada–Korea sovereign bond. *(published May 2026)*
WP-02	A Canada–United States Energy and Compute Compact. *(published May 2026)*
WP-03	A Canada–Korea Pacific Defence-Industrial Corridor. *(published May 2026)*
WP-04	The Addition Paradox — An energy thesis for Canada. *(published May 2026)*
TB-01	**The Verified Sky — Sensing, certainty, and the law of automated airspace awareness.** *(this briefing)*

AUTHOR	Jesse James
PUBLISHER	North Pacific Strategy Initiative
WEB	npsi.ca · canonical at npsi.ca/tb1/
CONTACT	linkedin.com/in/jessecares
LICENCE	Text CC-BY-4.0. The imprint and wordmark are not licensed.
CITATION	James, J. (2026). *The Verified Sky.* NPSI Technical Briefing No. 01.

END OF TECHNICAL BRIEFING NO. 01

Certainty is an architecture, not a model. The systems that work do not trust a classifier; they corroborate it, and they decline to act until the ledger closes.

— § II, CERTAINTY IS AN ARCHITECTURE

N O R T H P A C I F I C S T R A T E G Y I N I T I A T I V E

N P S I . C A