Ascension Venture Partners

Publication	Published	Cadence	Reading time
Public Edition v0.1	May 2026	Quarterly	~22 minutes

Maintained by

Ascension Venture Partners

The Warfront Taxonomy

A structured map of the technical sub-segments inside US defense autonomy and dual-use autonomous systems.

Ten sub-segments. Three tiers. Role types, comp band ranges, hiring patterns, talent clusters, and market structure observations across each. The structure is the asset — published because the depth becomes more valuable when it's recognized.

Download as PDF ↓

Read the taxonomy →

Public edition. Proprietary edition adds candidate benches, contact-level intelligence, and signal triggers.

---

00 Why this exists

The work happens at the level beneath the generic categories.

Most firms working in autonomous systems carry one or two generic categories — "aerospace," "defense," "robotics." That language is the giveaway.

The work happens at the level beneath: a maritime autonomy startup hiring its first head of perception is in a different talent market than an eVTOL company hiring its first head of certification engineering, even when both call themselves "autonomy."

The Warfront taxonomy is the structured version of that level beneath. We've been building it for over a year. It maps the technical sub-segments we work in, the role types that recur inside each one, the comp bands those roles command, the hiring patterns that show up at specific company stages, and the geographic clusters where the talent actually sits.

This is the public edition. It contains structural depth — what the segments are, who hires what, where the talent is. The proprietary edition we maintain internally adds candidate benches, contact-level intelligence, signal triggers, and prior placement history.

Founders and investors who want to see how the public taxonomy operationalizes against a specific role can scope a working session.

The structure is the asset. We publish the structure because the depth becomes more valuable when it's recognized.

01 How to read this

Ten sub-segments, organized in three tiers.

01–04

Tier 1

Vehicle classes

Where autonomy gets deployed. These segments organize around the physical platform.

05–07

Tier 2

Mission-layer software

What runs above the platform. These segments organize around the software stack that makes the platform useful.

08–10

Tier 3

Engineering disciplines

Where the talent actually specializes. These segments cut across vehicle classes — a perception engineer at a maritime startup and a perception engineer at a drone startup are in the same sub-segment of the talent market, even when their employers are not.

The three tiers are not exclusive. A typical Series A defense-autonomy company hires across all three. The taxonomy lets us match the specific shape of a hiring need to the specific shape of the talent pool — which is something that doesn't happen when a search firm carries "aerospace" as its only category.

Each entry below covers six things:

01 Definition — What the segment includes and excludes.

02 Role types — Recurring engineering and leadership roles.

03 Comp band ranges — What roles command in 2026 USD at US Series A through early B.

04 Hiring patterns — When companies hire what, and why.

05 Talent clusters — Where the people actually live and work.

06 Market structure — Capital flow, founder ecosystem, hiring rhythm.

Comp bands are presented as ranges. Exact numbers depend on stage, equity weighting, geographic location, and clearance status. The ranges below assume base salary at the median equity-tilted offer for a US Series A through early B autonomy company in 2026.

Tier 1

Sub-segments 01–04

Vehicle classes.

Where autonomy gets deployed. These segments organize around the physical platform.

01

Tier 1 · Vehicle classes

Aerial autonomy & UAS

01

Definition

Unmanned aerial vehicles operating across the size spectrum from sub-kilogram quadcopters to Group 5 fixed-wing platforms with wingspans exceeding 50 feet. Includes both fully autonomous platforms and operator-in-the-loop systems where autonomy reduces operator load. Excludes commercial passenger aviation (covered under eVTOL) and consumer drones below the prosumer line.

The segment splits operationally into three pools that draw on overlapping but non-identical talent: defense-mission UAS (ISR, strike, attritable platforms), commercial UAS (inspection, logistics, mapping), and counter-mission UAS (which sits in its own segment, see 05).

02

Role types

VP Engineering, Director of Engineering, Head of Flight Software, Head of GNC (guidance, navigation, control), Head of Avionics, Principal Flight Software Engineer, Principal Embedded Systems Engineer, Senior Staff Flight Test Engineer, Founding Engineer (early Series A), Director of Manufacturing & Production Engineering, Head of Mission Systems.

03

Comp band ranges

US, 2026

Role	Base	Equity
VP Engineering	**$280–400k**	+ 0.5–1.5%
Head of Flight Software / GNC	**$240–340k**	+ meaningful equity
Principal Flight Software / Embedded	**$230–320k**	+ equity
Senior Staff Flight Test	**$200–290k**	+ equity
Founding Engineer (Series A, pre–$10M ARR)	**$180–240k**	+ 1–3%

04

Hiring patterns

Companies in this segment hire flight software ahead of avionics, and avionics ahead of manufacturing. The first 25 engineering hires skew heavily toward flight software, GNC, and embedded systems. The shift

to manufacturing and production engineering happens at the Series B inflection, typically 12–18 months after Series A close. Defense-mission UAS companies add mission systems engineers earlier than commercial UAS companies because defense customers will not accept a platform without one.

05

Talent clusters

United States	Boston (MIT Lincoln Laboratory alumni, BAE Systems alumni), Bay Area (broad pool, highest comp pressure), San Diego (Northrop and General Atomics alumni — strong on Group 4–5 platforms), Austin (defense-tech wave, Anduril alumni), Seattle (Insitu alumni, Boeing).
United Kingdom	Bristol (BAE alumni, BRL spinouts), Cambridge.
Germany	Munich (Helsing-adjacent, MTU alumni).

06

Market structure

The segment compounds: Anduril alumni have founded the next generation of US drone companies, who are now hiring the next generation of operators, several of whom will themselves found drone companies in 24 months. The capital flow is reinforcing. Companies that don't hire ahead of the alumni-network curve get out-recruited by ones that do.

02

Tier 1 · Vehicle classes

Maritime autonomy

01

Definition

Uncrewed surface vessels (USVs) and uncrewed underwater vehicles (UUVs) operating in defense, scientific, and commercial maritime contexts. Includes both small-form-factor platforms (man-portable, ship-launched) and larger purpose-built autonomous vessels. Excludes manned vessels with assistive automation.

The defense subsection has matured fast since 2023 — Saronic's $4B valuation in 2024 marked a step change. Commercial maritime autonomy lags but is starting to follow.

02

Role types

VP Engineering, Director of Naval Architecture, Head of Marine Systems, Head of Autonomy (maritime-specific), Principal Subsea Engineer, Senior Staff Hydrodynamics Engineer, Director of Sea Trials & Operations, Head of Sensor Integration, Mission Systems Engineer (maritime-specific), Principal Communications Engineer (RF in maritime contexts).

03

Comp band ranges

US, 2026

Role	Base	Equity
VP Engineering	**$270–380k**	+ 0.5–1.5%
Director of Naval Architecture	**$230–320k**	+ equity
Head of Marine Systems	**$230–320k**	+ equity
Principal Subsea / Hydrodynamics	**$220–310k**	+ equity
Director of Sea Trials & Operations	**$200–280k**	+ equity

04

Hiring patterns

Maritime autonomy hiring is bottlenecked by naval architecture talent — there are far fewer trained naval architects than software engineers, and the discipline can't be cross-trained into from adjacent fields on a 12-month timescale. Companies that lock down naval architecture leadership early outperform companies that try to hire it on demand. Subsea hiring follows surface hiring by 6–12 months as platforms mature.

05

Talent clusters

United States	San Diego (Navy / SPAWAR alumni, defense maritime depth), Boston (subsea heritage, Bluefin alumni, Woods Hole adjacency), Hampton Roads / DC corridor (NSWC Dahlgren, Navy fleet adjacency), Austin (cross-recruited from defense-tech wave).
United Kingdom	Plymouth, Southampton (UK MoD adjacency, Thales naval).

Norway	Trondheim, Bergen (Kongsberg alumni — strong on subsea autonomy).
Sweden	Karlskrona (Saab Kockums alumni).

06

Market structure

This segment is the smallest of the Tier 1 vehicle classes by company count but among the highest by capital concentration. Defense maritime is structurally preferred by US DoD over commercial maritime for procurement reasons. Companies entering the segment from a defense-first frame outperform companies entering from a commercial-first frame on capital efficiency.

03

Tier 1 · Vehicle classes

Ground autonomy & defense vehicles

01

Definition

Uncrewed ground vehicles (UGVs) and autonomous-driving systems in both defense and dual-use commercial contexts. The segment includes light tactical UGVs, autonomous logistics vehicles (military and commercial), autonomous trucking, and platform-agnostic

autonomous-driving software stacks. Excludes consumer-vehicle ADAS (out of scope) and warehouse robotics (covered under industrial robotics, B-tier adjacent).

02

Role types

VP Engineering, Head of Autonomy Stack, Director of Vehicle Engineering, Head of Perception (vehicle-specific), Principal Motion Planning Engineer, Senior Staff Behavior Engineer, Director of Safety Engineering, Head of Hardware-in-the-Loop Test, Principal Sensor Integration Engineer, Head of Mapping & Localization.

03

Comp band ranges

US, 2026

Role	Base	Equity
VP Engineering	**$290–410k**	+ 0.5–1.5%
Head of Autonomy Stack	**$260–360k**	+ meaningful equity
Principal Motion Planning / Behavior	**$230–320k**	+ equity
Director of Safety Engineering	**$230–310k**	+ equity
Senior Staff Perception (vehicle)	**$220–300k**	+ equity

The autonomous-trucking subset commands a 10–15% premium on these bands due to talent compression.

04

Hiring patterns

Defense ground autonomy hires safety engineering later than commercial ground autonomy — DoD acceptance of "autonomy with operator override" is more permissive than civilian regulator

acceptance of full self-driving. Commercial autonomous trucking, by contrast, hires safety engineering and verification almost as early as core autonomy engineering. Companies that mismatch the hiring sequence to their actual customer (defense vs. civilian) burn capital on the wrong sequence.

05

Talent clusters

United States	Pittsburgh (CMU Robotics Institute, NREC, Argo alumni — the deepest single talent cluster in autonomous driving globally), Bay Area (Waymo, Cruise, Tesla alumni), Detroit (auto OEM alumni transitioning into autonomy startups), Boston (defense ground autonomy heritage, Endeavor Robotics alumni), Austin (defense-tech wave).
Germany	Munich (BMW Group autonomous-driving alumni, Helsing-adjacent ground), Stuttgart (Daimler alumni).
United Kingdom	Oxford (Oxbotica / Oxa alumni — strong commercial autonomy).

06

Market structure

The autonomous-driving segment went through a capital correction in 2023–2024 that pushed talent out of struggling companies and into adjacent segments — particularly defense ground autonomy and autonomous trucking. The talent quality available right now in 2026 is unusually strong relative to the funding available, which creates an arbitrage for well-capitalized Series A and B companies hiring at this moment.

04

Tier 1 · Vehicle classes

eVTOL & advanced air mobility

01

Definition

Electric vertical takeoff and landing aircraft for passenger and cargo applications, including urban air mobility (UAM), regional air mobility (RAM), and cargo-eVTOL platforms. The segment includes both clean-sheet eVTOL designs and electrified conversions of conventional rotorcraft. Excludes pure UAS (covered in 01) and traditional rotorcraft.

The segment is at the inflection between certification and commercial operation. The companies that close this transition successfully will define the segment for the next decade.

02

Role types

VP Engineering, Head of Certification, Director of Flight Test, Head of Powertrain Engineering, Principal Power Electronics Engineer, Senior Staff Aerodynamicist, Head of Avionics, Principal GNC Engineer, Director of Aircraft Systems, Head of Manufacturing Engineering.

03

Comp band ranges

US, 2026

Role	Base	Equity
VP Engineering	**$300–430k**	+ 0.5–1.5%

Role	Base	Equity
Head of Certification	**$260–370k**	+ meaningful equity (premium reflects scarcity)
Director of Flight Test	**$240–340k**	+ equity
Principal Power Electronics	**$230–320k**	+ equity
Senior Staff Aerodynamicist	**$220–310k**	+ equity

04

Hiring patterns

Certification is the bottleneck. The number of engineers in the world with eVTOL-relevant FAA Part 23 / EASA SC-VTOL certification experience is small enough to be effectively a known list. Companies that hire certification leadership early raise faster on subsequent rounds because investors price in lower regulatory risk. Companies that try to backfill certification at Series B regularly slip 12+ months on certification timelines.

Power electronics talent is the second bottleneck. The discipline cuts across automotive EV, eVTOL, and electrified marine — companies that don't compete on equity and mission against the EV sector lose the talent fight.

05

Talent clusters

United States	Bay Area (broad pool, Joby and Wisk adjacency), Boston (eVTOL heritage, MIT alumni), Wichita (Textron, Cessna conversion talent), Mojave / Edwards corridor (flight-test depth).

United Kingdom	Bristol (BAE, Vertical Aerospace alumni), Cranfield.
Germany	Munich (Lilium alumni — strong but distributed post-restructuring), Bavaria more broadly.
Switzerland	Zurich (ETH-trained electrified-flight talent), Bern.
Austria	Vienna and Linz (CycloTech alumni, AVL alumni — strong on power systems).

06

Market structure

The eVTOL segment is unusually concentrated by capital. Five companies (Joby, Archer, Volocopter, Lilium successor entities, Vertical Aerospace) account for the majority of disclosed equity capital raised in the segment. Series A and B companies in eVTOL are competing for talent against companies that have raised 10–50× more. This means the smaller companies need to compete on something other than comp — typically mission specificity, technical interest, or equity weighting. Recruitment in this segment that ignores this dynamic underperforms.

Tier 2

Sub-segments 05–07

Mission-layer software.

What runs above the platform. These segments organize around the software stack that makes the platform useful.

05

Tier 2 · Mission-layer software

Counter-UAS & active defense systems

01

Definition

Detection, tracking, identification, and neutralization systems for hostile uncrewed aerial threats. The segment includes RF/radar/EO sensor systems, mission software for threat assessment, and effector platforms (kinetic, electronic warfare, directed energy). Excludes general air defense (covered under traditional defense primes).

The segment has compressed dramatically since 2023 as drone use in active conflict has shifted procurement priorities. Capital and procurement attention have both tilted hard toward US-based counter-UAS startups in the last 18 months.

02

Role types

VP Engineering, Head of Sensor Engineering, Head of Mission Software, Principal Radar Engineer, Principal RF Engineer, Senior Staff Computer Vision Engineer (drone-detection-specific), Director of Threat Assessment Algorithms, Head of Effector Integration, Principal Electronic Warfare Engineer.

03

Comp band ranges

US, 2026

Role	Base	Equity
VP Engineering	**$290–410k**	+ 0.5–1.5%
Head of Sensor / RF Engineering	**$250–350k**	+ meaningful equity
Principal Radar Engineer	**$240–340k**	+ equity (premium reflects scarcity)
Principal RF Engineer	**$230–320k**	+ equity
Director of Threat Assessment	**$230–310k**	+ equity

04

Hiring patterns

Radar engineering is the structural bottleneck. The talent pool was built around a small number of legacy defense primes and has not been replenished at the rate required by the current wave of counter-UAS startups. Companies that hire radar leadership early can attract the rest of the team around that anchor; companies that don't, can't.

Mission software hiring follows sensor engineering by 6–9 months. Companies that try to build the mission software layer ahead of the

sensor layer overspecify the software against assumptions about sensor capabilities that turn out to be wrong.

05

Talent clusters

United States	DC corridor / Northern Virginia (defense radar heritage, Raytheon alumni, MITRE adjacency), Boston (defense radar, MIT Lincoln Laboratory), San Diego (defense electronics, Northrop alumni), Huntsville (Army Missile Defense adjacency), Austin (defense-tech wave).
United Kingdom	Edinburgh (Leonardo radar heritage), Bristol.
Other	Israel-trained engineers relocated to the US are a meaningful population in this segment, though geographically distributed.

06

Market structure

The segment is the most clearance-sensitive of the Tier 2 categories. Roughly 60–70% of senior engineering roles require active or eligible US security clearance. AVP places cleared-eligible candidates in this segment; cleared-only roles are deferred to specialized cleared-search partners.

06

Tier 2 · Mission-layer software

Mission software & command-and-control

01

Definition

The software layer that turns autonomous platforms into operational capabilities — mission planning, command-and-control (C2), data fusion across multiple platforms, operator interfaces, and the AI/ML systems that surface decisions to operators. Includes both startup-built mission software (Anduril Lattice, Shield AI Hivemind) and the broader category of platform-agnostic mission tools.

Mission software is where defense autonomy companies are increasingly differentiated. The platform layer has become commoditized faster than the mission software layer, and capital is following the differentiation.

02

Role types

VP Engineering, Head of Mission Software, Director of Platform Engineering, Principal Distributed Systems Engineer, Senior Staff Backend Engineer (with defense data domain experience), Head of UI/UX (defense-context-specific), Principal AI/ML Engineer (decision systems), Head of Data Engineering, Director of Integration Engineering.

03

Comp band ranges

US, 2026

Role	Base	Equity
VP Engineering	**$300–430k**	+ 0.5–1.5%
Head of Mission Software	**$270–380k**	+ meaningful equity
Principal Distributed Systems	**$250–350k**	+ equity
Principal AI/ML (decision systems)	**$260–370k**	+ equity
Head of UI/UX (defense context)	**$230–310k**	+ equity

04

Hiring patterns

Mission software companies hire backend and distributed systems engineers earlier than platform-engineering or UI/UX. The first 30 engineering hires are typically 70%+ backend / distributed systems. UI/UX hiring tends to lag, and the lag is a recurring failure mode — defense customers are notoriously sensitive to operator-interface quality, and companies that under-invest in UI/UX get returned by user testing late in development cycles.

05

Talent clusters

United States	Bay Area (Palantir alumni, ex-FAANG with defense interest, Anduril alumni), Austin (defense-tech wave), DC corridor (defense software heritage, Booz Allen alumni transitioning to startups), Boston (defense software, MITRE adjacency), Seattle (cross-recruited from Microsoft, Amazon defense practices), Pittsburgh (CMU software talent transitioning to defense).

United Kingdom	London (defense software, Faculty AI alumni).
Germany	Munich (Helsing alumni — strong on AI/ML decision systems).

06

Market structure

The Anduril effect compounds in this segment more than in any other Tier 2 category. The Lattice platform has redefined what defense customers expect from mission software, which raises the floor for every other entrant. Companies entering the segment now have to clear a quality bar that was lower 24 months ago. This compresses the talent market — the same engineers who could have been first-50 hires at Anduril in 2022 are being recruited to start companies in the same segment in 2026.

07

Tier 2 · Mission-layer software

Sensor fusion & perception stack

01

Definition

The full sensor pipeline from raw sensor input through perception output: lidar, radar, camera, GNSS-denied positioning, multi-sensor

fusion, scene understanding, object detection and tracking. The segment cuts across vehicle classes — perception engineers serve aerial, maritime, ground, and space autonomy contexts, and the underlying skills transfer cleanly between them, though deployment specifics differ.

02

Role types

VP Engineering (perception-led), Head of Perception, Director of Sensor Engineering, Principal Computer Vision Engineer, Principal Sensor Fusion Engineer, Senior Staff Lidar Engineer, Principal Radar Algorithms Engineer (perception-side), Head of Calibration & Sensor Models, Principal SLAM Engineer, Head of Scene Understanding.

03

Comp band ranges

US, 2026

Role	Base	Equity
VP Engineering (perception-led)	**$290–420k**	+ 0.5–1.5%
Head of Perception	**$260–370k**	+ meaningful equity
Principal Computer Vision	**$240–340k**	+ equity
Principal Sensor Fusion	**$250–360k**	+ equity (premium reflects scarcity)
Principal SLAM	**$230–330k**	+ equity

04

Hiring patterns

Perception is the most cross-segment-fluid talent pool in autonomous systems. A senior perception engineer working on autonomous trucking can move to maritime autonomy or eVTOL with 3–6 months of

ramp. This makes the segment unusually competitive — perception leadership at any single company is constantly being recruited by every other autonomy company hiring perception leadership. Retention is structurally hard.

The corollary: perception bench depth is the single strongest recruitment moat in autonomous systems. Companies that maintain pre-qualified perception benches across vehicle classes can fill perception roles in weeks where competitors take quarters.

05

Talent clusters

United States	Pittsburgh (CMU Robotics Institute — the deepest perception cluster globally), Bay Area (Tesla, Cruise, Waymo perception alumni), Boston (MIT, BU computer vision), Seattle (Microsoft Research, autonomous-systems perception), San Diego (defense perception, Qualcomm autonomous adjacency).
United Kingdom	Oxford (Oxford Robotics Institute — global-tier on SLAM and perception), Cambridge, Edinburgh.
Germany	Munich (BMW autonomous perception alumni, Helsing-adjacent), Stuttgart.
Switzerland	Zurich (ETH Computer Vision Lab — globally ranked), Lausanne (EPFL).

06

Market structure

This segment has the strongest cross-vehicle-class transfer of any Tier 2 or Tier 3 category. The implication for hiring: a recruitment partner that organizes its bench by vehicle class loses against one that organizes by discipline. AVP organizes perception by discipline first, then segments by application context — which is why perception roles fill faster against our bench than against vehicle-class-organized search firms.

Tier 3 Sub-segments 08–10

Engineering disciplines.

Where the talent actually specializes. These segments cut across vehicle classes — a perception engineer at a maritime startup and a perception engineer at a drone startup are in the same sub-segment of the talent market, even when their employers are not.

08

Tier 3 · Engineering disciplines

Autonomy stack: GNC, motion planning, behavior

01

Definition

The decision-making layer of autonomous systems: guidance, navigation, and control (GNC) for aerospace platforms; motion planning and behavior generation for ground and maritime platforms; high-level mission planning and autonomy arbitration. The segment is the connective tissue between perception (what the system sees) and actuation (what the system does).

02

Role types

Head of Autonomy Stack, Principal GNC Engineer, Principal Motion Planning Engineer, Senior Staff Behavior Engineer, Principal Control Theory Engineer, Senior Staff Reinforcement Learning Engineer (autonomy-applied), Director of Autonomy Architecture, Head of Decision Systems.

03

Comp band ranges

US, 2026

Role	Base	Equity
Head of Autonomy Stack	**$260–370k**	+ meaningful equity
Principal GNC Engineer	**$230–330k**	+ equity (premium for aerospace-cleared)
Principal Motion Planning	**$230–320k**	+ equity
Senior Staff Behavior Engineer	**$220–310k**	+ equity
Senior Staff Reinforcement Learning (autonomy)	**$240–340k**	+ equity

04

Hiring patterns

Autonomy stack hiring is bottlenecked by control-theory talent at the senior and principal levels. Strong control engineers are not interchangeable with strong ML engineers, despite a category of confusion that emerged in 2023–2024. Companies that hire ML engineers and ask them to do control engineering ship buggy systems; companies that hire control engineers ahead of the curve outperform.

The architecture of the autonomy stack tends to ossify around the first principal hires. Companies that hire well in the first 5–7 autonomy stack roles benefit for years; companies that hire poorly carry the architectural debt for years. This is one of the highest-leverage hiring sequences in autonomous systems.

05 Talent clusters

United States	Pittsburgh (CMU — NREC and Robotics Institute), Bay Area (Stanford, Berkeley control theory alumni), Boston (MIT Aerospace Controls Laboratory), Atlanta (Georgia Tech).
United Kingdom	Oxford (ORI), Cambridge (Engineering Department), Imperial College London.
Switzerland	ETH Zurich (Institute for Dynamic Systems and Control — globally ranked).
Germany	Stuttgart, TU Munich.

06 Market structure

The senior end of this segment is one of the smallest globally — true principal-level GNC and control-theory engineers number in low thousands worldwide, not tens of thousands. Companies that maintain warm relationships with this population over multi-year horizons (rather than only contacting them when a role opens) get the early signal when senior people consider moves. AVP's bench in this segment is built on multi-year relationship cultivation, not transactional outreach.

09

Tier 3 · Engineering disciplines

Edge ML & embedded inference

01

Definition

Machine learning systems that run on resource-constrained hardware in autonomous platforms — onboard inference for perception, decision, and control loops, model compression and quantization, custom inference accelerators, and the integration of ML into real-time embedded systems. Excludes pure cloud-side ML (out of segment) and consumer edge ML (out of segment).

02

Role types

Head of Edge ML, Principal Embedded ML Engineer, Senior Staff Model Optimization Engineer, Principal Hardware-Software Integration Engineer (ML focus), Director of Inference Infrastructure, Principal ML

Compiler Engineer, Senior Staff Deep Learning Engineer (embedded systems experience).

03

Comp band ranges

US, 2026

Role	Base	Equity
Head of Edge ML	**$260–360k**	+ meaningful equity
Principal Embedded ML	**$240–340k**	+ equity
Senior Staff Model Optimization	**$230–320k**	+ equity
Principal ML Compiler	**$250–360k**	+ equity (premium reflects scarcity)
Director of Inference Infrastructure	**$230–310k**	+ equity

04

Hiring patterns

Edge ML hiring is constrained by the intersection of two relatively scarce skill sets — strong ML engineers who also have deep embedded systems experience. The Venn diagram is small. Companies that try to hire each side separately and integrate them through process tend to ship slower than companies that hire engineers who span both. The latter cost more, hire harder, and ship faster.

This segment is where the talent scarcity is most acute relative to demand right now. Capital flow into autonomous platforms outpaces edge ML talent supply by a wide margin. The result: edge ML hiring is taking 6–9 months for senior roles even at well-funded Series A and B companies.

05

Talent clusters

United States	Bay Area (Apple Silicon, Tesla, Cruise embedded ML alumni), Seattle (Microsoft Research embedded ML, Amazon hardware ML), Boston (MIT, embedded ML research), Pittsburgh (CMU embedded systems), Austin (Apple, embedded compute).
United Kingdom	Cambridge (ARM alumni — strong intersection of ML + embedded compute), Bristol (Graphcore alumni).
Switzerland	Zurich (ETH integrated systems), Lausanne.
Germany	Munich (Infineon, Bosch automotive ML).

06

Market structure

The hardware accelerator generational shift (Jetson Thor, custom silicon, NPU integration) has created a 2026-specific dynamic: engineers with experience on the previous generation of hardware are hireable but require 3–6 month ramp to be productive on current platforms. Companies that screen on prior platform experience without considering ramp are filtering out a meaningful percentage of the actually-good talent pool. AVP's screen in this segment looks at ramp trajectory, not historical platform mix.

10

Tier 3 · Engineering disciplines

Simulation, test, and safety case engineering

01

Definition

The full lifecycle of autonomous-system verification and validation: high-fidelity simulation environments, hardware-in-the-loop test, scenario coverage analysis, safety case construction, formal verification, and certification engineering for both aviation (FAA, EASA) and ground-vehicle (NHTSA, UNECE) regulators. The segment cuts across vehicle classes but is most weight-bearing in eVTOL and autonomous trucking, where regulatory acceptance is the bottleneck to commercial deployment.

02

Role types

Head of Verification & Validation, Director of Safety Engineering, Principal Simulation Engineer, Senior Staff Hardware-in-the-Loop Test Engineer, Head of Certification, Principal Safety Case Engineer, Director of Test Operations, Senior Staff Formal Verification Engineer, Principal Coverage Analysis Engineer.

03

Comp band ranges

US, 2026

Role	Base	Equity
Head of V&V	**$250–360k**	+ meaningful equity
Director of Safety Engineering	**$230–320k**	+ equity (premium for cleared)
Head of Certification	**$260–370k**	+ meaningful equity (acute scarcity)
Principal Simulation Engineer	**$220–310k**	+ equity
Principal Safety Case Engineer	**$240–340k**	+ equity
Principal Formal Verification	**$250–360k**	+ equity

04

Hiring patterns

This is the most under-staffed engineering function in autonomous systems at Series A and B. Founders prioritize hiring the building functions (perception, autonomy stack, vehicle engineering) and defer V&V hiring until regulatory pressure forces it. The pattern repeats across companies and across segments. Companies that buck the pattern and hire V&V leadership in the first 30 engineering hires consistently move through certification and customer acceptance faster than companies that defer.

Certification engineering specifically is the single hardest-to-fill role in autonomous systems. The number of engineers globally who have led certification of an autonomous platform from start to entry-into-service is in the low hundreds. Companies that recruit for this role transactionally — opening a JD when they need it — typically wait 9–14 months. Companies that maintain pre-qualified certification benches fill in 8–14 weeks.

05

Talent clusters

United States	Wichita (FAA-corridor certification heritage, Cessna and Beechcraft alumni), Mojave / Edwards corridor (flight-test depth), Bay Area (Joby, Wisk certification alumni), Boston (MIT, MITRE simulation), Pittsburgh (CMU, NREC test engineering), Detroit (auto V&V transitioning into autonomous trucking).
United Kingdom	Cranfield (certification heritage), Bristol (BAE certification).
Germany	Stuttgart (auto V&V alumni — strong on autonomous trucking transition), Munich.
France	Toulouse (Airbus certification alumni).

06

Market structure

This segment is where AVP's specialism is most differentiated against generalist autonomy recruitment. Most search firms in autonomous systems do not maintain dedicated certification or V&V benches — the talent pool is too small to support generalist sourcing, and the role specifications are too specialized to fill from broad search. The depth here is built across years and doesn't replicate quickly. It's also where founders most often realize, late, that they should have hired earlier.

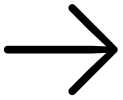

Bridge · publication-as-credibility to publication-as-CTA

How AVP works against this taxonomy.

The taxonomy is the organizing principle of how AVP operates. Concretely:

01

Bench depth maps to sub-segment, not to vehicle class.

AVP maintains pre-qualified specialist benches organized by the engineering disciplines in Tier 3 (perception, autonomy stack, edge ML, V&V) rather than by the vehicle classes in Tier 1. The talent transfers across vehicle classes more than it transfers across disciplines. A company hiring a Head of Perception for a maritime platform recruits from the same bench that a company hiring a Head of Perception for a drone platform does — the candidates are fluent in the discipline; the deployment context is the second filter.

02

Roles get scoped against three Tier-aligned signals.

When AVP takes a brief, the first three questions filter the search against the taxonomy. Which vehicle class? — sets the deployment context. Which mission-layer function? — sets the architectural surface. Which engineering discipline? — sets the bench to draw from.

The three together produce a tight specification that most generalist briefs miss.

03

Sub-segment fluency is the entry point on every engagement.

AVP does not take engagements where the prospect cannot articulate which sub-segment of the taxonomy their hiring need lives in. The conversation happens during scoping, not after. Companies that haven't thought through where their need sits in this structure are not yet ready for the kind of targeted recruitment AVP runs — and we'd rather not take an engagement than run one against a poorly-specified brief.

04

Geographic depth maps to talent clusters, not to office locations.

AVP's geography is wherever the talent for a given sub-segment actually clusters. We don't have offices in Pittsburgh or Zurich — we have multi-year relationship investment in those clusters because that's where the perception and autonomy stack benches live.

Proprietary edition

What's behind the wall.

This is the public edition of the taxonomy. The proprietary edition we maintain internally adds, for each sub-segment, the operational layer that turns the structural map into a usable bench.

01 Specific named candidates currently on the bench, with current employment status and motivation indicators.

02 Contact-level intelligence per company in the sub-segment — who decides, who blocks, who's been there how long, what the hiring thesis is.

03 Exact compensation data: real offer histories, real counter-offer dynamics, real equity packages by stage.

04 Signal pattern triggers: what events, in what sequence, predict which companies are about to hire which roles.

05 Prior placement history per sub-segment, where confidentiality permits.

If the structure of the public edition fits a need you're trying to solve — scope a working session.

Not a pitch. A working session means we sit down against your specific role, against the sub-segment it lives in, and walk through how the proprietary bench operationalizes against it. Please specify the sub-segment your need lives in when you reach out — generic enquiries can't be scoped properly.

Start a conversation — a.rennie@ascension-international.com

Ascension Venture Partners

San Diego · Boston · DC

Edited by Alex Rennie

v1.0 — May 2026 · US defense autonomy & dual-use autonomous systems