MANIFESTO FUCK HUMANOIDS. INVEST IN HEAVY ROBOTICS.

GENOMA LABS · FUSE PLATFORM

The autonomy layer for *heavy robotics.*

FUSE is the open integration + proprietary control stack that lets the world's largest hydraulic machines perceive, decide, and act under autonomous command. Strategic OEM anchor: CPS Group. Patent filed. Validated on four machines.

ENTITY	ROUND	OEM ANCHOR	IP POSTURE
Genoma Labs, Los Angeles CA	Seed · $5M	CPS Group · 20,000+ units	USPTO Provisional Filed

§ 02 · CATEGORY THESIS · WHY NOW

CATEGORY HEAVY ROBOTICS INCUMBENTS NONE

Heavy machines still have no lattice.

Cars got autonomy stacks. Drones got autonomy stacks. The world's largest hydraulic robots still operate through fragmented controls, manual workflows, and machine-by-machine integration.

01 - STILL MANUAL

Operator-bound workflows

Construction, lifting, marine, and utility workflows remain dependent on scarce human operators in hazardous environments.

02 - STILL FRAGMENTED

Locked inside OEM stacks

Control systems, telemetry, and autonomy experiments remain trapped inside OEM-specific stacks and bespoke deployments.

03 - STILL OPEN

No common software layer

No common software layer yet coordinates sensing, policy execution, and machine control across hydraulic robotic fleets.

§ 03 · PLATFORM · THE OPEN AUTONOMY LAYER

FUSE is the autonomy layer.

FUSE does for hydraulic robotics what modern autonomy platforms did for drones and defense systems: connects machines, sensors, policies, and operators inside one deployable layer.

DEPLOYMENT SURFACE	Hydraulic robotic fleets
INTEGRATION POSTURE	Open below · proprietary above
COMPOUNDING	Policy + perception + fleet data

01 - OPEN BELOW

Integrates, not replaces

FUSE integrates with machine controls, autonomy stacks, and external robotics software instead of requiring a closed full-stack rewrite.

02 - INTELLIGENT MIDDLE

One runtime for large robots

Normalizes sensing, control, tasking, and execution into one runtime for large hydraulic robotic systems.

03 - EXPANDS ABOVE

Compounds with every deployment

Every deployment adds reusable policy intelligence, perception services, and cross-machine autonomy workflows.

§ 04 · ARCHITECTURE · OPEN INTEGRATION / PROPRIETARY INTELLIGENCE

Open below.
Proprietary above.

Open where the market must integrate. Proprietary where Genoma Labs can own policy execution, perception, and fleet intelligence - the layers that compound with every deployment.

PROPRIETARY · ABOVE				COMPOUNDS WITH SCALE
POLICY & INTELLIGENCE	POLICY MODELS **Distilled Policy Models** Domain models trained on proprietary ROS execution policies for hydraulic tasks.	EXECUTION INTELLIGENCE **Intent → Action Runtime** Translates intent into robust machine actions under real control, load, and safety constraints.	SENSOR AI LAYER **Decentralized Perception** Edge nodes fusing local sensing, machine state, and environment awareness.	FLEET ORCHESTRATION **Cross-machine Services** Coordinates behaviors, monitoring, updates, and reusable autonomy workflows.
OPEN · BELOW				INTEROPERABILITY LAYER
INTEGRATIONS	MACHINE CONTROLS **OEM interfaces** Hydraulic control systems, telemetry buses, and lawful machine-access pathways.	ROS / AUTONOMY STACKS **External runtimes** Mission runtimes and policy executors needing a practical bridge into heavy machines.	SENSORS & PAYLOADS **Multi-modal inputs** Cameras, lidar, inertial systems, task sensors, and machine-specific payloads.	THIRD-PARTY SYSTEMS **Vessels, vehicles, remote ops** Partner autonomy products that plug into the FUSE layer.

§ 05 · SENSE · MACHINE PERCEPTION & AWARENESS

PERCEPTION · THE SENSING HALF

SENSE perceives. FUSE acts.

A six-modality perception layer mounted directly on heavy machinery: LiDAR, 3D cameras, radar, sonar, ultrasonic and laser ranging fused into a single all-weather, day-night, through-object picture of the machine's surroundings.

01 **All-weather, day-night perception**
Works where LiDAR fails - dust, fog, rain, darkness.

02 **Through-object awareness**
Radar sees beyond walls, vehicles, and dust plumes.

03 **Six-modality fusion, one stack**
Single on-machine unit → single source of truth for FUSE.

SIX-MODALITY STACK

M01 LiDAR	M02 3D Cameras	M03 Radar
M04 Sonar	M05 Ultrasonic	M06 Laser Range

DATAFLOW · SENSE → FUSE

SENSE — 6-modality fusion → World Model — Continuous 3D scene → FUSE — Policy & actuation

§ 06 · SCOPE · DISCIPLINE · WHAT SENSE IS / WHAT IT IS NOT

SCOPE DISCIPLINE · FOR THE AVOIDANCE OF DOUBT

What SENSE is. What it is not.

A clean boundary is a compounding asset. SENSE is a native on-machine perception stack paired with FUSE - not a bolt-on, not a sensor, not a vision-only startup. Scope discipline separates a platform from a feature.

WHAT SENSE IS ✓ IN SCOPE	WHAT SENSE IS NOT ✕ OUT OF SCOPE
01 **Native on-machine perception stack** Mounted directly on heavy machinery - hardware, firmware, ML.	01 ~~An aftermarket bolt-on~~ Retrofit kits don't capture the machine data flywheel.
02 **Multi-modal · six sensors fused** LiDAR · 3D cameras · radar · sonar · ultrasonic · laser ranging.	02 ~~A LiDAR-only system~~ Single-modality stacks break in dust, fog, darkness, occlusion.
03 **Tier-1 OEM integrated** Designed into the machine, not retrofitted after the fact.	03 ~~A consumer product~~ Industrial-grade only - priced and warranted to fleet owners.
04 **Safety-gated by default** ISO 15143 / 18497 / EN 474 pathways - operator-protective.	04 ~~A vision-only startup~~ Cameras are one of six modalities - not the product.
05 **Paired with FUSE autonomy** Perception → world model → policy → machine - a closed loop.	05 ~~A standalone sensor~~ SENSE without FUSE is half the loop - we ship both.

A platform, not a peripheral. Scope is the moat.

§ SCOPE DISCIPLINE

§ 07 · MODEL LAYER · FUSE + SENSE · VALIDATED AUTONOMY

Sub-decimeter autonomous control. Validated on four machines.

FUSE is the autonomy controller that drives hydraulic machinery to precision targets - trained in simulation, deployed on-machine. **SENSE** is the perception layer that feeds it. Together they close the loop from world-awareness to actuator command.

END-EFFECTOR PRECISION
0.096 m
Validated in closed-loop simulation, reproduced on real hardware.

INDEPENDENT MACHINES
4
Distinct hydraulic platforms validated end-to-end under FUSE control.

RECORDED TRAINING RUNS
850+
Proprietary corpus of hydraulic trajectories, actions, and outcomes.

01 - POLICY SUBSTRATE

Structured execution policies

Domain-specific reinforcement learning trained on hydraulic valve dynamics. The policy learns curriculum-constrained joint expansion - starting simple, earning complexity - until it reaches full 16-DoF precision control (12 crane + 4 outrigger). 28-DoF on roadmap.

02 - DISTILLED MODELS

Long-term edge

GPU architecture is a first-class training variable. Checkpoints are hardware-distribution-bound - identical code on a different GPU produces a different model. The trained stack is not a commodity anyone can reproduce by reading the patent.

03 - RUNTIME ACTION LAYER

Precision control policy

The trained policy runs on an edge device at real-time rates. The edge device holds actuator authority - the server requests, it never commands. Safety enforcement is local, not cloud-dependent.

04 - COMPOUNDING FEEDBACK

Every deployment improves the stack

Each machine running SENSE + FUSE generates labeled multi-modal data that tightens the next model. Deployed fleet → better perception → better policy → faster deployment.

§ 08 · FUSE · THE AUTONOMY CONTROLLER

SENSE PERCEIVES · FUSE ACTS

The decision & action half of the autonomy pair.

FUSE is a reinforcement-learning controller that drives articulated hydraulic machinery to **sub-decimeter precision** without human input. It runs a trained policy on an edge device, enforces real-time safety constraints over CAN bus, and learns entirely inside a physics simulation before its first contact with the real machine.

PRECISION	MACHINES	RUNS	TO VALIDATION
0.096 M	4 ×	850 +	<30 DAYS

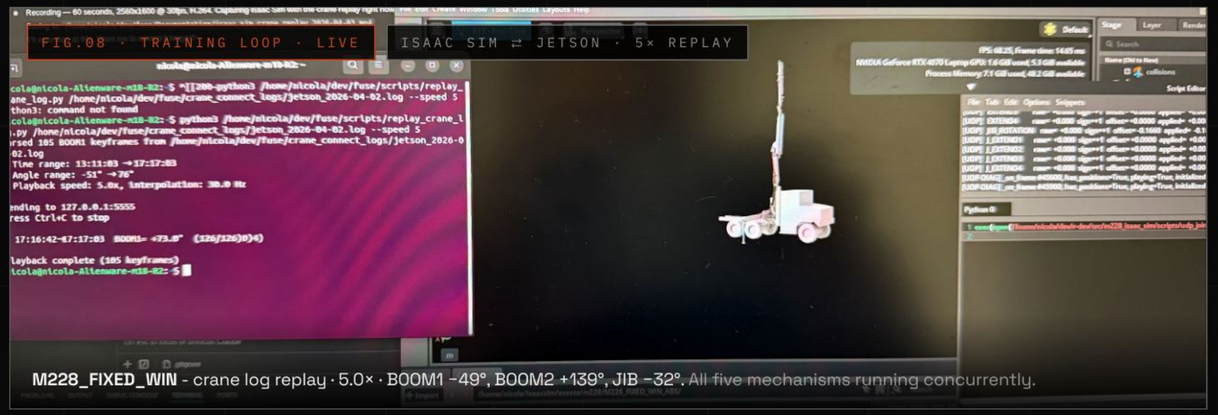

M228_FIXED_WIN - crane log replay · 5.0× · BOOM1 −49°, BOOM2 +139°, JIB −32°. All five mechanisms running concurrently.

01 - CURRICULUM-CONSTRAINED DOF EXPANSION

Stepwise from 4 → 16-DoF

Joint set expands in controlled increments under an empirically-validated max-expansion constraint. Prevents the entropy saturation that otherwise collapses training at high DoF - the mechanism that makes sim-to-real tractable on complex hydraulics.

02 - ALGORITHM TRANSITION GATE

0.87 m → 0.096 m · *9× precision jump*

A convergence monitor detects when on-policy methods hit a reproducible ceiling across seeds, then gates a transition to off-policy for the final precision phase. Fires on evidence of ceiling, not a fixed schedule.

03 - RECURSIVE PRECISION REFINEMENT

≥30% distributional overlap per hop

Target distribution narrows under a mathematically-derived overlap constraint - not tuned. Narrowing faster causes immediate regression; the constraint prevents catastrophic forgetting while precision compounds.

04 - HARDWARE-AWARE DISTRIBUTED TRAINING

GPU architecture = first-class hyperparameter

Identical code, seed & hyperparameters diverge across GPU microarchitectures. Cross-GPU checkpoint transfer degrades precision 1.7–2.8× (validated). Same-GPU continuation improves 0.92×. Checkpoints are not portable commodities.

05 - DIGITAL TWIN CALIBRATION PIPELINE

Machine-readable JSON, architecture-neutral

Calibration maps physical sensor readings to simulation joint states, resolves sign conventions, and emits a JSON instruction set. Any inference runtime can reproduce the sim-to-real mapping without access to the calibration process.

§ 09 · SCIENTIFIC METHOD · OEM-GRADE DATA → PHYSICAL LAW → DETERMINISM

OEM-grade data. Not startup-grade. That's the moat.

Physical AI at this size of robotic machines can only be built by an OEM. The data required is not startup-grade: not spotty demos, not borrowed sessions, not synthetic proxies. It is **OEM-grade**: proven physical systems run continuously, with full-stack telemetry, under warranty, for years. We own the fleet, we own the sensors, we own the CAN bus, we own the build floor. We don't need permission to collect physical-world data. **We just run the machines, like we always have.**

AXIOMS OF THE METHOD

01 OEM-grade data, not startup-grade. Continuous runtime on proven systems under warranty. Our fleet, our CAN bus, our sensors. Spotty demo sessions cannot model physical law. Only OEMs can.

02 The probabilistic space is the medium, not the limitation. Physics is probabilistic at small scale, deterministic at the scale of cranes, hydraulics, and electromagnetic waves.

03 A pattern is real only when it is reproducible across independent measurements: seeds, GPU architectures, model scales. Anything else is a single draw.

04 The fleet is the experiment. Three seeds in parallel, cross-architecture validation, cross-model-scale ablations. Run until the pattern is unmistakable.

05 Reward floors are messages from the physical world. The −0.010 SAC floor pointed to missing physics in the reward. We listened, re-shaped, and the fleet now climbs to R = +0.358.

PATTERNS ALREADY FOUND · PHYSICALLY REPRODUCIBLE

LORA RANK — **Capacity, not data volume, is the bottleneck.** 0.072 ±10%

r = 16 → 32 → 64 → 128 → 256 yields loss 1.16 → 0.37 → 0.20 → 0.105 → 0.072. Three independent GPU architectures (Turing / Ampere / Ada) converge to the same r=256 floor within ±10%.

SAC BREAKTHROUGH — **We listened to the floor. The floor broke.** R = +0.358

R = −0.010 was a message: missing physics in the reward. After reward re-shaping the fleet climbed off the floor. **s37 (16-DoF)** now at R = **+0.351** @ 11.3M steps. On pace to challenge s999's +0.566 peak. **s1 (28-DoF)** steady at R = **+0.358**.

SENSE CSI — **Phase-slope regression is electromagnetic law.** 1 cm / 488ch

488+ WiFi subcarriers → 1 cm spatial accuracy. Not a neural approximation of presence detection: it **is** wave propagation, expressed as computation. TX diversity scaled 1 → 4 today. (SENSE runs the LAMSA algorithm family internally.)

ENSEMBLE DISTILL — **Quality of hypothesis space beats quantity.** −40% loss

Ensemble data (3× pairs) at r=128: **0.1756** vs **0.1054** alone. Reproducible across every rank tested. Corpus pipeline now feeds 1.7B + 27B simultaneously: our first cross-scale data point.

GPU ARCHITECTURES
3 × 3
Turing · Ampere · Ada. Cross-architecture attractor

SAC · FLEET PEAK · LIVE
R = +0.358
s1 28-DoF @ 4.8M steps · s37 16-DoF climbing at +0.351 / 11.3M steps.

WIFI SUBCARRIERS · CSI
488+
Phase-slope → sub-centimeter spatial resolution

DATA PROVENANCE
OEM-grade
CPS Group fleet. Proven systems, continuous telemetry, under warranty.

§ 10 · LIVE TRAINING LOG · 0.094 M PRECISION · SUCCESS THRESHOLD CROSSED

0.094 m positional precision. Success threshold crossed.

Success = reach target within 0.100 m (100 mm). The v5-reward / UTD=0 run on **GS 3080 · seed 13** hit **0.094 m (94 mm)** in a single best episode: the fleet record. Three other runs tied at **0.098 m (98 mm)**: AW, NW, NA. **Sub-0.1 m crossed.** Reward (R) is the search signal; precision is the physical outcome. Physics just said yes.

SEARCH SIGNAL · SAC REWARD · T+0

● LIVE

MACHINE	GPU	JOB	STEPS	FPS	R	NOTES
GS	5070	s53 · 16-DoF	3.2M / 30M	867	+0.093	Warming up
GS	3080	s53 · 16-DoF	12.9M / 30M	881	+0.181	Steady
GS	4070S	s37 · 16-DoF	11.3M / 30M	**910**	+0.351	Strong. Challenging s999's +0.566 peak
AW	TITAN	s1 · 28-DoF	4.8M / 30M	505	+0.358	Stable. Fleet peak

NEXT STEP. PUSH PRECISION FURTHER // HIGH R **AND** TIGHT TARGETING, TOGETHER

01 Take the best CORRECTFLAGS checkpoint **(s999, R = +0.566)**

02 Fine-tune with tight targeting `--target-r-min 0.02 --target-r-max 0.15 --reward-version 5`

03 Sub-0.1 m precision emerges from high reward **plus** tight targeting

PRECISION · FLEET RECORD

0.094 m

GS 3080 · seed 13 · single best episode · v5 reward / UTD=0

TIED AT

0.098 m

AW · NW · NA. Three independent runs, tight-range era

SUCCESS THRESHOLD

< 0.100 m

Industrial-safety spec. Now crossed by 4 runs

NEXT CAMPAIGN

s999 → fine-tune

R=+0.566 checkpoint + target-r 0.02 to 0.15. High R **and** tight targeting

§ 11 · PROOF OF LIFE · REAL STEEL · SIM-TO-REAL LOOP CLOSED

NOT A RENDER · NOT A DECK MOCK

A COPMA 228 autonomous crane, mounted on an MTVR, driven by FUSE.

A 16-DoF articulated hydraulic system - 12 crane joints plus 4 outrigger legs - calibrated to an Isaac Sim digital twin on a USMC Medium Tactical Vehicle Replacement. The filed sim-to-real pipeline, running today against real steel.

Cutting-edge OEM hydraulic technology on a military-grade wheeled chassis - already a NATO supplier.

The lawful path from the CPS Group commercial fleet (20,000+ units, 50+ countries) into defense-grade autonomous logistics and engineering support.

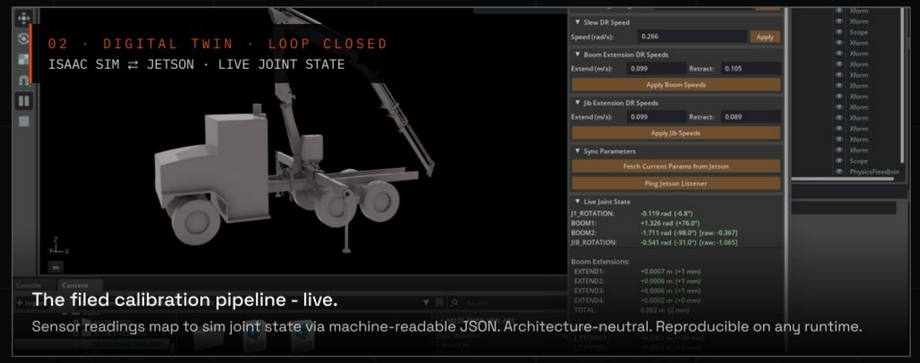

The filed calibration pipeline - live.

Sensor readings map to sim joint state via machine-readable JSON. Architecture-neutral. Reproducible on any runtime.

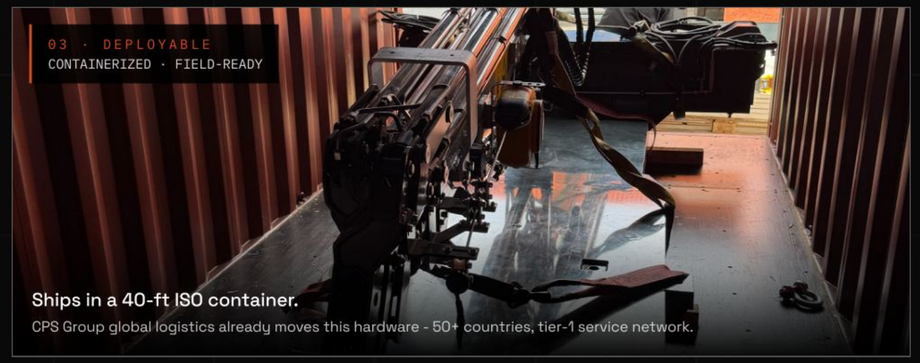

Ships in a 40-ft ISO container.

CPS Group global logistics already moves this hardware - 50+ countries, tier-1 service network.

CRANE	PLATFORM	DoF	BRIDGE	STATUS
COPMA 228	USMC MTVR	16 · crane + outrigger	Jetson ⇄ Isaac Sim · UDP	Running today

§ 12 · FUSE · VALIDATION SCOPE & ROAD TO SERIES A

Filed, reproduced, running on real steel.

What's on the record in Provisional #590734290 · and the 90-day priorities that take this from seed proof to Series-A-ready defense-tech.

VALIDATION SCOPE · AS FILED	PROVISIONAL #590734290
MACHINES **4** independent hydraulic cranes	GPU TRAINING NODES **6** heterogeneous architectures
RECORDED RUNS **850+** full provenance	SEEDS · PRECISION **4/4** sub-decimeter reproduced
INTEGRATION **CAN-bus on production OEM hardware**	EDGE AUTHORITY **Jetson-class · actuator-local**

SAFETY ARCHITECTURE · TWO-BUS DESIGN

NEXT TWO QUARTERS · SERIES A PRIORITIES

G·01 Ops hours on real steel
Machine-hours, site count, near-miss reductions, operator-override incidents. Even 10 hours at a live CPS Group site beats any simulation claim. *Film it.*

G·02 Operator override rate - *declining*
A single number: how often a human overrides the autonomous controller. Lower is better; a declining trendline across weeks is the story.

G·03 Precision vs. human baseline
File 0.096 m against a measured skilled-operator precision figure on the same crane. If FUSE is tighter than the operator, that is the deck headline.

G·04 Cycle time · lifts per hour
Operators care about throughput, not meters. Any cycle-time reduction versus a human-driven baseline becomes the economic headline.

G·05 Safety architecture · *strongest differentiator*
Two-bus design: edge holds actuator authority, server only requests. Industrial & defense investors recognise this as correct architecture - most autonomy startups get it wrong.

Filed innovations, reproduced precision, correct safety posture. Gaps are execution, not invention.

§ 13 · SENSOR AI LAYER · DISTRIBUTED RUNTIME

Decentralized perception. Unified runtime.

FUSE distributes sensing and execution across machine-level nodes, then coordinates them through a common runtime and shared policy layer - local enough for safety-critical execution, unified enough to become infrastructure.

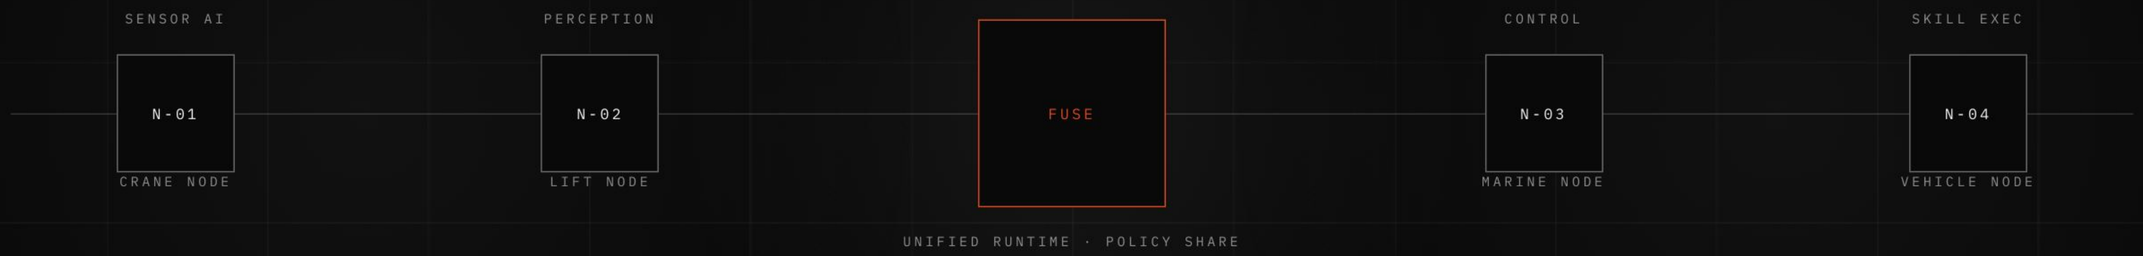

PERCEPTION AT THE EDGE

Sensor AI nodes process camera, telemetry, and machine-state inputs locally - each machine understands its environment in real time.

CONTROL WHERE IT MATTERS

Execution policies run close to the hydraulic machine interface, enabling faster response, safer behavior, and tighter closed-loop control.

LEARNING ACROSS THE FLEET

Each node participates in a broader platform that shares models, updates policies, and improves autonomy across machines over time.

§ 14 · GO-TO-MARKET · WEDGE

WEDGE · ARTICULATED CRANES

Start with cranes. Expand everywhere.

Articulated cranes are the first deployment surface: hard hydraulics, real economic pain, direct access to high-value industrial workflows. They force FUSE to solve the real problem early.

MACHINE CLASS	Articulated hydraulic cranes
CONTROL CHALLENGE	16-DoF nonlinear hydraulics
PARTNER ACCESS	CPS Group OEM pathway
EXPANSION SURFACE	Vessels · utility · defense · adjacent fleets

The crane wedge is the beachhead, not the ceiling.

§ 15 · OEM ANCHOR · CPS GROUP · UNFAIR ADVANTAGE

OEM ANCHOR · CPS GROUP · COPMA / PESCI MARINE

Tier-1 OEM distribution at founding.

Most autonomy startups spend years chasing their first OEM integration. FUSE starts with one: CPS Group - COPMA articulated cranes, PESCI Marine, and the broader CPS platform - shipping into our target fleets today.

UNITS SHIPPED

20,000+

CPS Group machines in the field - the beachhead fleet for FUSE integration.

COUNTRIES

50+

Global service & distribution footprint, pre-built dealer network.

ANNUAL SALES

€25M

Proven commercial base - FUSE rides an existing revenue stream.

MULTIPLE · STANDALONE

~2×

What Tier-1 OEM distribution is worth to a defense-tech investor.

01 · COMPOUNDING ACCESS

Every CPS unit is a FUSE seat

The installed base becomes the distribution channel. No cold-start on OEM integrations - FUSE ships where CPS already ships.

02 · HARDWARE TRUTH

Real 16-DoF hydraulics

COPMA cranes force the autonomy layer to solve articulated hydraulic control on day one - not a simulated toy problem.

03 · DUAL-USE LANE

PESCI Marine → defense vessels

CPS marine products open a lawful path from civilian hydraulics into USV / UGV defense programs.

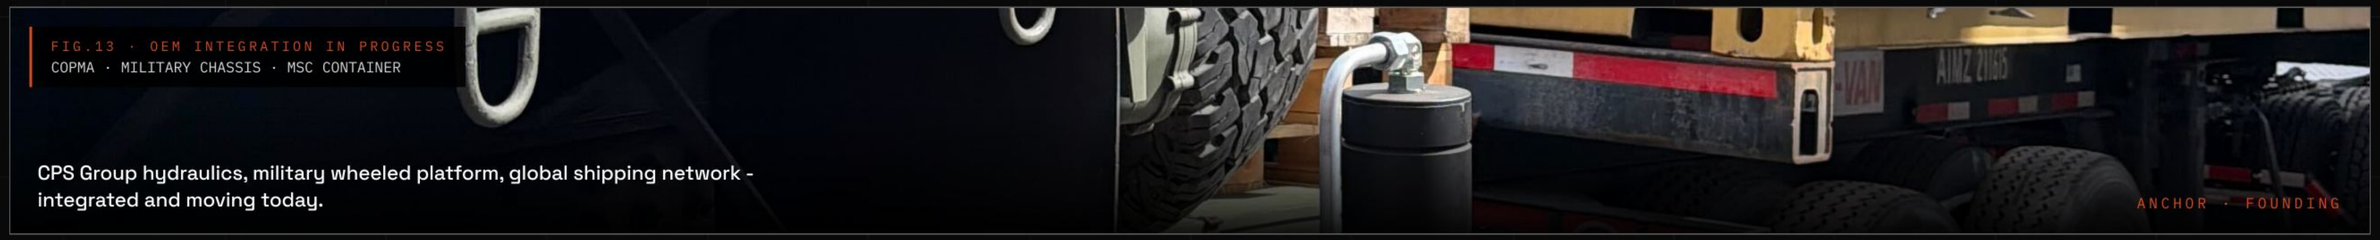

FIG.13 · OEM INTEGRATION IN PROGRESS
COPMA · MILITARY CHASSIS · MSC CONTAINER

CPS Group hydraulics, military wheeled platform, global shipping network - integrated and moving today.

ANCHOR · FOUNDING

Competitors need to win OEM deals. We started with one.

§ 16 · MARKET · TOTAL ADDRESSABLE OPPORTUNITY

MARKET · TAM / SAM / SOM

A trillion-dollar installed base of hydraulic work.

FUSE addresses the entire installed base of heavy articulated machines - cranes, vessels, ground vehicles, and industrial robotics - each running unamortized labor today and priced on a per-machine autonomy layer tomorrow.

Every articulated hydraulic machine is a seat for an autonomy layer.

TAM · TOTAL ADDRESSABLE

$1.2 T

Global heavy-machine autonomy surface · cranes, USVs, UGVs, mining, forestry, ports, construction.

SRC: OFF-HIGHWAY RESEARCH 2024 · MCKINSEY HEAVY EQUIP 2023

SAM · SERVICEABLE

$180 B

Articulated hydraulic fleets in OECD + EU/UA defense operators reachable via OEM integration.

SRC: ALLIED MARKET RESEARCH 2024 · SIPRI DEFENSE 2024

SOM · 2028 BEACHHEAD

$2.4 B

Cranes + dual-use maritime via CPS & partner OEMs, 5-yr capture.

SRC: CPS GROUP INSTALLED BASE · INTERNAL BOTTOM-UP

INSTALLED UNITS

7.4 M

Heavy articulated hydraulic machines, global fleet.

LABOR DISPLACED

$420 B/YR

Annual operator hours addressable across the installed base.

OEM PARTNERS · TARGET

40 +

Integrations to cover 80% of SAM by 2030.

UNIT ECONOMICS · PER MACHINE	INTEGRATION (ONE-TIME)	FUSE LICENSE (ARR)	DATA / MODEL UPSIDE	GROSS MARGIN	CUSTOMER PAYBACK
SaaS + integration, attach rate rises as CPS ships.	**$18 K** SENSE harness + commissioning. ~55% GM.	**$24 K/YR** Autonomy layer subscription. Software.	**+$6 K/YR** Fleet data / policy-as-a-service to OEMs.	**~82 %** Blended at steady state. Software-led.	**< 9 MO** vs $180K/yr operator cost displaced.

§ 17 · REVENUE MODEL · EVERY MACHINE A RECURRING NODE

Every deployed machine becomes a recurring node.

BLENDED PER MACHINE

~$30K / YR

Autonomy layer + data services. Hardware sale on entry.
~82% gross margin at steady state.

01 · HARDWARE SALE

ONE-TIME · ENTRY POINT

SENSE harness + FUSE compute module

$15–18K

PER MACHINE INSTALLED

- Sensor harness + integration
- Commissioning & calibration
- First install via CPS fleets

GM ~55% · FUNDS THE FLYWHEEL

02 · FUSE LICENSE

RECURRING · THE BIG LINE

Autonomy subscription per machine

$24K / YR

PER MACHINE / ARR

- Autonomy stack runtime
- OTA policy updates
- Fleet orchestration + monitoring

SOFTWARE GM · ~90%

03 · MARKETPLACE

RECURRING · COMPOUNDING

Take rate on third-party autonomy apps

20%

OF APP REVENUE ON FUSE

- Autonomy SKILLS marketplace
- Task packs from integrators
- Digital twin library

THE LATTICE FLYWHEEL

04 · DATA & MODELS

RECURRING · DEFENSIBILITY

Policy-as-a-service back to OEMs

+$6K / YR

PER MACHINE / ATTACH

- Fleet telemetry licensing
- Distilled policies to OEMs
- Compute / sim credits

DATA MOAT · ROADMAP

VALUE DISPLACED

$180K / YR per operator

What $30K/yr of FUSE replaces on the machine.

CUSTOMER PAYBACK

< 9 MO

Hardware + first-year FUSE recovered inside a fiscal year.

ATTACH · CPS CHANNEL

Factory-fit

Sold with the crane, not retrofit. Attach rises with every ship date.

§ 18 · FLEET SCALE · SOFTWARE REVENUE SCALES WITH FLEET SIZE

Three units of fleet, three units of ARR.

The path from pilot to platform is linear in machine count, multiplicative in software revenue. Each machine is a recurring node: hardware on entry, FUSE license every year, marketplace & data on top.

FLEET REALITY CHECK

US knuckle-boom crane sales alone run ~3,000 units / year. Reaching a 3K–5K deployed fleet globally is a capture problem, not a TAM problem.

BLENDED ARPU	STEADY-STATE GM	NET DOLLAR RETENTION
~$30 K / YR	~82 %	> 120 %

Software revenue scales with fleet size - the moat compounds one machine at a time.

§ 19 · FIELD EVIDENCE · UKRAINE IS WINNING WITH ROBOTS

Ukraine is winning because of field robotics.

Trucks go autonomous. Vessels go autonomous. **The lifting machinery sitting on them must be autonomous too — and must talk to the host platform's autonomy stack.** That handshake is FUSE. **Robots learn even when destroyed** — joint states, maneuvers, and telemetry relay back. People don't.

01 · THE HANDSHAKE

Autonomous truck + autonomous crane = one control loop. The host platform's stack and the lifting machine's stack must speak the same language. **FUSE is that language.**

02 · ROBOTS LEARN. PEOPLE DON'T.

A destroyed robot relays joint states, forces, perception frames, and maneuver history back to the fleet. Every loss is a training sample. Human attrition has no such upside.

03 · MULTI-OEM OR NOTHING

Ukraine doesn't procure from one prime — it buys from dozens of OEMs and wires them together. **The winning architecture is horizontal, not vertically-integrated.**

04 · HYDRAULICS ARE NOT OPTIONAL

Logistics, engineering, mine-clearing, recovery — the war runs on hydraulic actuation. Wheels drive. Booms lift. Arms dig. All of it gets autonomy, or none of it does.

§ 20 · DEFENSE FLAGSHIP · FROM DESIGN TO REALITY

Next one: 95 metric tons of lift - we are not showing, we are executing.

Building the **COPMA 950** (95 T/m knuckle-boom) onto the **Oshkosh M1070 HET** - the US Army's heaviest tactical truck. First-of-its-kind autonomous heavy loader, teleop + autonomous via FUSE + SENSE.

TARGET · OCT 2026

FIRST OF ITS KIND

NOW IN BUILD · LOS ANGELES

FROM DESIGN → REALITY

VISION (RENDERS) · BUILD IN PROGRESS (PHOTOS, LA YARD)

PHASE 00 · RENDER

DESIGN
M1070 x COPMA 950 - flagship vision

PHASE 01 · RENDER

PATTERN VALIDATED
MTVR + COPMA - proven integration pattern

PHASE 02 · REAL · LA YARD

PARTS · BOTH CHASSIS
COPMA hardware in yard, trucks staged

PHASE 03 · REAL · LA YARD

INTEGRATION
M1070 chassis on stands, crane mounting

PHASE 04 · REAL · LA YARD

SIDE BY SIDE
MTVR (validated) + M1070 (flagship)

01 · LIFT CAPACITY

95 T/m COPMA 950 knuckle boom - heaviest hydraulic loader ever integrated onto a tactical HET chassis.

02 · DUAL MODE

Teleoperation from the truck cabin or over satellite, plus full autonomous execution - same FUSE, same SENSE, same policies.

03 · PLATFORM

Oshkosh M1070 HET is the US Army's primary tank transporter - a validated procurement path with a live global fleet.

04 · WARP SPEED

Flagship in build now, target Oct 2026 - from commitment to first-of-kind autonomous HET in under 12 months.

EXPORT · DUAL-USE POSTURE			
	EAR 9x515 / ECCN-aware architecture. Commercial dual-use core; defense variants gated behind license workflow.	**CAGE-code supplier** (CPS Group) · **EU NIS2** governance body (ACN, Apr 2026). Regulated on both sides.	**ITAR-clean by design** — civilian-first, defense adjacencies reached through authorized partners only.

§ 21 · EXECUTION · YEARS OF WORK · DTLA · CAD → SIM → STEEL

Code is now accessible. Hardox and Weldox Steel are not.

We run a full factory: design integration, digital twin, controls, fab, RL and deployment (real-world data).

One roof, one team: **OEM digital twin access with CAD (Rhino, SolidEdge, Onshape etc)**, **RL policy development on a live 7-DoF ABB 2400 and 6400** and **in-house welding & integration**. The loop from simulation to deployment closes fast. **Years of work in DTLA, not Hawthorne. Not cheap to build here. That is the moat.**

NO VENDORS · NO GATE-KEEPERS

CAD → SIM → STEEL

DTLA · NOT HAWTHORNE

YEARS OF WORK · CAPITAL-HEAVY

THE OPERATION · NOT A RENDER

01 LEARN · 02 FAB · 03 FLOOR · 04 DESIGN · 05 CONTROLS

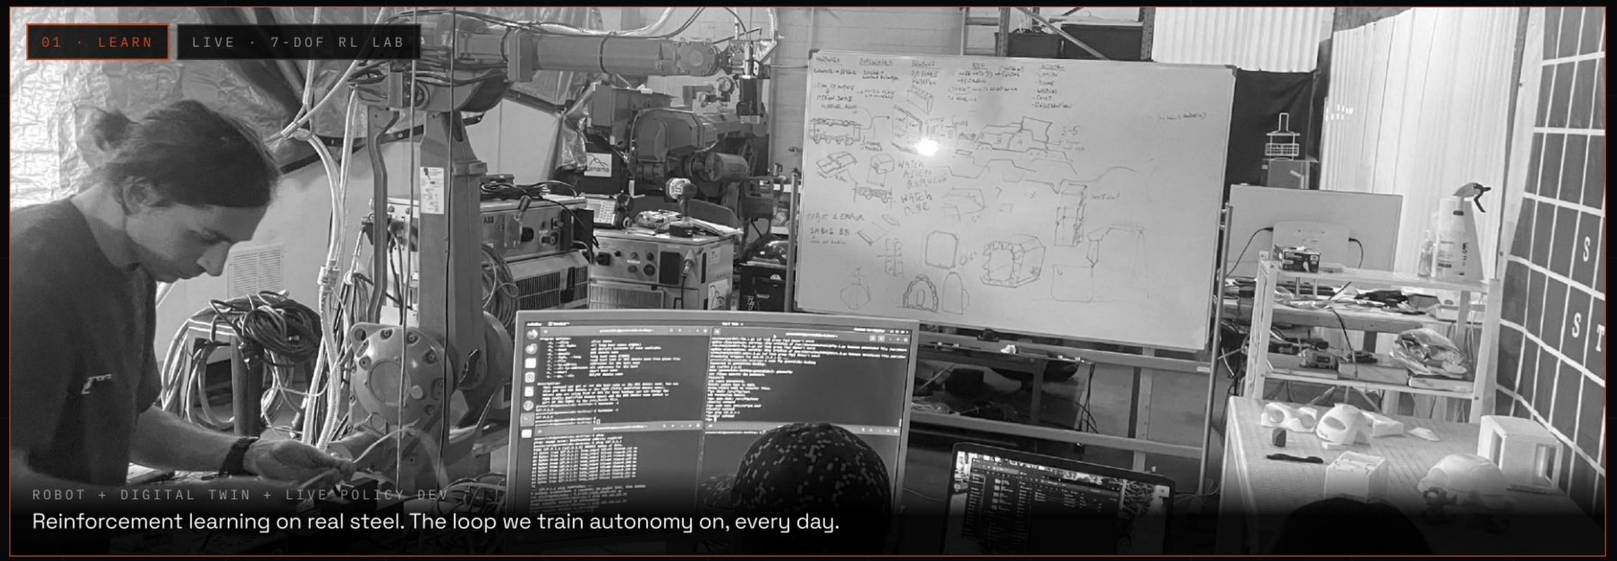

01 · LEARN

LIVE · 7-DOF RL LAB

ROBOT + DIGITAL TWIN + LIVE POLICY DEV

Reinforcement learning on real steel. The loop we train autonomy on, every day.

02 · FAB

WELD · MOUNT

We cut, weld, and integrate our own hardware. No subcontractor critical path.

03 · FLOOR

DUAL ROBOT · ENVIRO BOOTH

Purpose-built LA shop. Training rigs, enviro testing, and integration under one roof.

04 · DESIGN

DIGITAL TWIN · RHINO · SOLIDEDGE

CAD · COPMA + CHASSIS FIT

Every bracket modeled before it is cut. Twin drives sim, sim drives build.

05 · CONTROLS

WIRING · SENSORS · TEST RIGS

Every harness and ECU bench-tested before it touches the machine.

§ 22 · COMPETITIVE · HEAD-TO-HEAD · M-MET PROGRAM

COMPETITIVE · M-MET PROGRAM · U.S. ARMY

We outfield the trade-show concept.

AM General, Carnegie Robotics, and Textron Systems announced a collaborative agreement to design an autonomous M-MET - a next-generation unmanned ground vehicle for the U.S. Army. It is a trade-show concept; FUSE + SENSE is a fielded autonomy layer for exactly this machine class.

CAPABILITY	AM GENERAL · CARNEGIE · TEXTRON	GENOMA · FUSE + SENSE
Autonomy model class	Classical motion-planning stack, vendor-integrated	VLA + RL policy, distilled on real hydraulic traces
Perception	LiDAR + camera, line-of-sight only	Six-modality SENSE - through-object, all-weather, day-night
Hydraulic actuation (crane + boom)	Not demonstrated on booth unit	16-DoF articulated hydraulics - CPS Group OEM pathway
Field posture	Trade-show concept, Oct 2025 announcement	Deployable platform across civilian & dual-use fleets
Integration posture	Three-party consortium, single program vehicle	Open below, proprietary above - platform, not a vehicle

A consortium built one concept truck. We built the autonomy layer every heavy machine - M-MET included - will need to run on.

§ 23 · DEFENSIBILITY · COMPOUNDING ADVANTAGE

The moat compounds with deployment.

FUSE becomes more defensible every time it touches a real machine, executes a policy, and learns from an operating fleet. The moat is not one model - it is a growing control, policy, and deployment graph.

01 · MACHINE ACCESS

OEM pathways

Partner-enabled OEM pathways and control integrations are hard to replicate and grow more valuable as FUSE spans more machine types.

02 · EXECUTION POLICIES

Proprietary ROS policies

Policies capture how large hydraulic robots should move, react, and complete work under real machine constraints.

03 · DEPLOYMENT DATA

Field telemetry loop

Failure modes, operator behavior, and perception feedback create a training loop improving models and runtime decisions.

04 · REUSABLE WORKFLOWS

Validated autonomy

Each workflow redeploys across adjacent machines, sites, and autonomy stacks instead of starting from zero.

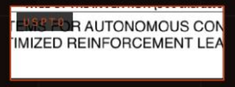

PROVISIONAL · USPTO · FUSE

Methods and Systems for Autonomous Control of Hydraulic Machinery via Optimized Reinforcement Learning

FILED BY FOUNDER · ASSIGNED TO CPS GROUP S.P.A. (PROPRIETARY MACHINERY IP & DIGITAL TWIN)

APP #590734290

SENSE PROVISIONAL · FILING NOW

Second USPTO filing in flight - machine-perception stack. **Window closes on inbound claims as the filing is docketed.**

NVIDIA — INCEPTION MEMBER

GPU credits, Jetson track, technical review.

DEEPMIND — ROBOTICS RESEARCH · CANDIDATE

Under consideration by Google DeepMind robotics.

WHY NOT THE INCUMBENTS?	NVIDIA	CAT · KOMATSU (OEMS)	BOSCH · PARKER (TIER-1)
Their business model doesn't let them.	Isaac/Cosmos are horizontal sim stacks, not vertical field autonomy. **We are their customer** — Inception member, distilled model runs on Jetson.	Cat Trimble + Komatsu FrontRunner are single-OEM stacks, 15+ years in, **still don't cross-integrate**. Can't become the horizontal layer.	Component vendors. Bosch Rexroth ships BODAS; **BODAS is a PLC**, not a policy stack. No autonomy muscle, no OEM-graph.

The moat is not one model or one machine.
It is a growing control, policy, and deployment graph that gets stronger as the platform scales.

§ 24 · OPERATING MODEL · SKUNK WORKS, AGENTIC

OPERATING MODEL · NOT A STARTUP

We don't run like a startup. We run like Skunk Works.

Kelly Johnson built the SR-71 Blackbird - **the fastest aircraft ever flown, still unmatched 60 years later** - with fewer than 75 engineers. Not because Lockheed had unlimited resources. Because Johnson's team had radical autonomy, proprietary tooling, and zero organizational drag.

Genoma Labs runs the same model. **The enabling instrument is Matrix** - our proprietary agentic AI platform. A 255-capability autonomous engineering system closing the full loop: research, code, RL training across 6 GPU nodes, sensor-corpus generation, patent drafting, CAN-bus hardware integration, multi-machine orchestration - all coordinated agents.

> A competitor with **200 engineers in silos** cannot ship faster than a **Matrix-augmented team in a single agentic loop.** The SR-71 was never outrun by a larger program. Neither will this one be.

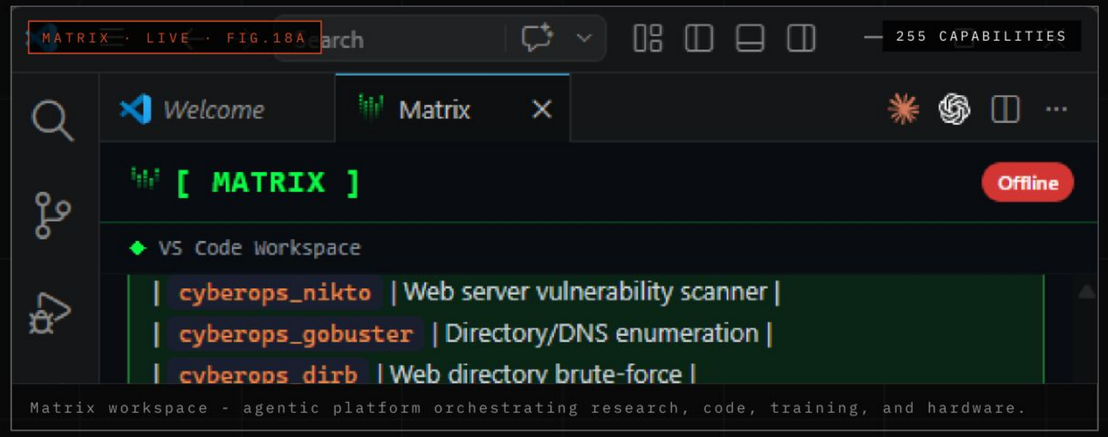

Matrix workspace - agentic platform orchestrating research, code, training, and hardware.

SKUNK WORKS · 1962
SR-71 Blackbird
< 75 ENGINEERS · MACH 3.3
UNMATCHED · 60 YEARS

GENOMA · 2025
FUSE + SENSE + Matrix
SMALL CORE TEAM · 850+ RUNS
2 PATENTS · 6-MODALITY

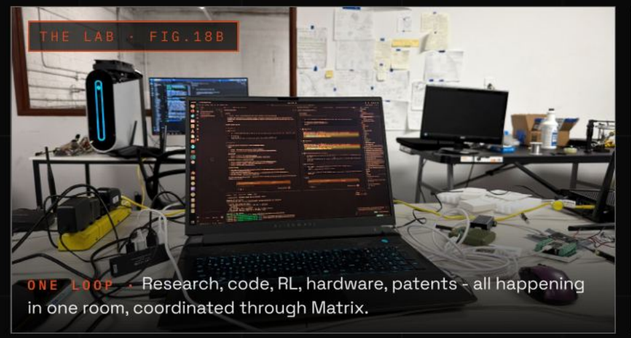

ONE LOOP · Research, code, RL, hardware, patents - all happening in one room, coordinated through Matrix.

MATRIX · CLOSES THE LOOP

01 Real-time research
02 Code generation
03 RL training · 6 GPUs
04 Sensor corpus gen
05 Patent drafting
06 CAN-bus integration
07 Multi-machine orchestration
08 Distillation pipeline

ALL OUTPUTS BUILT THROUGH MATRIX.

ROUND POSTURE · SPECIAL INVITE

A Seed round, for defense and robotics / edge funds who live in this problem space. Not easy. Not a moonshot. Inevitable. And we are best-positioned to win it. Technology validated, patent filed, OEM channel live across **50+ countries**, proprietary model stack in training. A **limited, special-invite round** for investors who understand what it means when a company this technically deep reaches this stage this early.

WINDOW · NARROW BY DESIGN

§ 25 · DETERMINISTIC VLA · PHYSICAL-AI CONTROL AT HYDRAULIC SCALE

A 100% deterministic VLA. Does not exist at this scale.

Vision-Language-Action **is** the control policy · and we are building the first **100% deterministic VLA** for heavy hydraulic machines. RL grounds the policy on real steel; distilled LLMs are the substrate VLAs run on, sized for our edge silicon. **16-DoF hydraulic control, sub-decimeter precision, four independent machines.** No VLA at this scale and this determinism exists today · not at OpenAI, not at DeepMind, not at Figure. Determinism is the investor-grade property: same inputs, same action, every time · audit-ready, safety-certifiable, contract-ready.

VS. OEM INCUMBENTS

Trimble, Hexagon, Caterpillar own sensor integration - but have no RL training infrastructure and no LLM strategy.

VS. AUTONOMY STARTUPS

Startups can build one layer - perception, control, or a foundation model. They cannot close the loop across all three without the OEM anchor we have and they don't.

THE LOCK

Checkpoints are hardware-distribution-bound - they cannot be reproduced by reading our patents or copying our code. Every machine we deploy trains the next model.

LIVE TRAINING MATRIX · FIG.18

MODEL	TYPE	STATUS	TRAINING DATA	GPU · TARGET
- ROBOTIC MOTION -				
SAC s999	RL · 16-DoF	● 15.3M · R=+0.34	15.3M env steps	RTX 3080 → Jetson
SAC s100 (NA)	RL · 16-DoF	● 17.4M · R=+0.19	17.4M env steps	RTX 4070L → Jetson
SAC s200	RL · 16-DoF	✓ 30M · R=+0.20	30M env steps	TITAN (done)
qwen3-14b-vla-v1	VLA · LoRA r=64	✓ Merged · GGUF	6,841 pairs	TITAN (build)
SENSE RL1 n=3	RL · channel-sel	✓ 10-seed valid	10,400 fusion windows	RTX 4070L
- CODING & SERVER LLMS -				
gemma4-31b-codex-v1	QLoRA r=64 NF4	● Corpus 4792/5000	6,841 + ~5,000 traces	AW TITAN
matrix-fuse-v1	QLoRA deployed	✓ Live	1,424 genoma-brain pairs	All machines · Ollama
jackrong-27b	Q4_K_M distill	✓ Live (offline)	Pre-trained (HF)	RTX 3090 · Ollama

✓ TRAINED / LIVE ● IN TRAINING HETEROGENEOUS GPU FLEET · RTX 3080 → TITAN

The moat is not the algorithm. The moat is the closed loop between real steel, real data, and a model stack that only gets stronger the more machines run it.

§ 26 · FOUNDING TEAM · OPERATORS, NOT TOURISTS

Not a first-time founder.
A 25-year operator.

Genoma Labs is led by a 25-year operator embedded in the OEM anchor. **Formally designated to CPS Group S.p.A.'s administrative & executive body under EU NIS2** (Italian National Cybersecurity Agency, Apr 2026) · advisor & manager of CPS US, the US subsidiary · **CAGE-code-qualified defense supplier**. Decades across heavy equipment engineering, hydraulic manufacturing, cybersecurity, and marketing. Not a Stanford grad learning heavy industry for the first time.

FOUNDER · CEO

NICOLA MONTUSCHI

25 years across heavy equipment: engineering, manufacturing, cybersecurity, and marketing. **Designated member of CPS Group S.p.A.'s administrative & executive body under EU NIS2** (ACN, Apr 2026). Advisor & manager of **CPS US** · CAGE-code-qualified defense supplier. Filed the USPTO provisional. Built the FUSE controller and SENSE stack personally.

CPS GROUP S.P.A. · NIS2 GOVERNANCE BODY CPS US · ADVISOR / MANAGER

HYDRAULIC SYSTEMS OEM MANUFACTURING CYBERSECURITY NVIDIA INCEPTION

OEM ANCHOR · REGULATED SEAT

NIS2 governance body
CPS Group S.p.A.

+ advisor / manager, CPS US

EU NIS2 administrative body at the parent (ACN filing, Apr 2026) · CAGE-code defense supplier at US sub. COPMA, PESCI Marine, 50+ countries. Family-held. First-name access to engineering, ops, and commercial leadership.

ADVISOR · AUTONOMY / CONTROL

Luca Montuschi

Control System Group Leader · Visa Cash App RB F1 Team

F1-grade real-time control, sensor fusion, deterministic edge loops — the exact stack FUSE runs on.

ADVISOR · MACHINE INTEGRATION

Pierluigi De Maria

Electrical Mgr, CPS Group · ex-EFFER / HIAB

Legacy integration, CAN/LIN ↔ electro-hydraulic distribution. Decades inside the machines we retrofit.

HIRING · NEXT 18 MONTHS **6** Robotics / RL **4** Firmware / Edge **3** Field Ops **2** OEM Integration 15 HIRES · FUNDED BY ROUND

§ FOUNDING TEAM

Decades inside the machine. This is not a first attempt.

§ 27 · CATEGORY · VERTICAL AUTONOMY PLATFORMS · OUTCOMES

Vertical autonomy platforms compound into $10B+ outcomes.

When a single company owns the autonomy layer for a critical vertical - defense, intel, industrial data - the outcome is not a margin business. It is a category-defining one. Hydraulic work is a larger, earlier, and more fragmented TAM than any of these comparables.

DEFENSE AUTONOMY	INTEL DATA LAYER	MARITIME AUTONOMY	HYDRAULIC AUTONOMY
Anduril	**Palantir**	**Saildrone**	**Genoma Labs · FUSE**
$28B VALUATION · 2024	$300B+ MARKET CAP	$1.2B · FROM SEED	$1.2T TAM · OPEN CATEGORY
Owned the software layer above U.S. defense platforms. Vertical integration from sensor to policy.	Became the default data / decision layer for intelligence and government. Platform, not product.	Autonomous USV fleet. Built the vertical autonomy stack for maritime surveillance · now a dual-use defense platform.	Autonomy layer for every hydraulic machine on Earth. No incumbent. Patent filed. OEM anchored.

LARGER

$1.2T of installed hydraulic work globally - larger than defense + automotive combined.

EARLIER

No Anduril, no Palantir, no Tesla-FSD equivalent exists for hydraulic machinery. **Category is open.**

ANCHORED

CPS Group OEM distribution from day one. Not a cold-start integration play.

Anduril owned defense. Palantir owned intel. Genoma Labs will own hydraulics.

SEED · THE ASK

$5M

To take FUSE from four validated machines to a deployed fleet under CPS Group distribution.

ROUND Seed - Priced	TERM 18 months to Series A
LEAD [LEAD INVESTOR]	PRICING Valuation inline with dual-use robotics seeds

USE OF FUNDS · 18 MONTHS

40%	**Engineering · R&D** Robotics, RL, edge runtime. 10 senior hires.	$2.0M
28%	**Field deployment · OEM integration** CPS Group pilot fleet. 50+ units under FUSE control.	$1.4M
20%	**Training infrastructure** GPU fleet, sim-to-real pipeline, data ops.	$1.0M
12%	**G&A · IP · Legal** Patent conversion, export controls, dual-use compliance.	$0.6M

MILESTONES TO SERIES A

- **50+** machines deployed under FUSE · CPS-distributed
- **3** dual-use contracts · civilian + defense adjacencies
- **Sub-5cm** field-validated precision · operator-outperforming
- **$3-5M** booked ARR + $2M integration revenue · blended $80-100K / machine

§ 29 · MACHINE MOAT · THE FLYWHEEL

The machine flywheel compounds.

From the machine perspective, the moat is a self-reinforcing loop: proprietary data, simulation fidelity, real-world execution, and skill distillation - tied to real hydraulic systems rather than generic robotics abstractions.

GPU HOURS
1,200+
ML modeling on hydraulic behavior.

DOF MODELED
16+
Crane + outrigger. 28 on roadmap.

SIM PRECISION
<10 CM
Isaac Sim sim-to-real.

FEEDBACK LOOP
N→N+1
Distillation every deployment.

01 · PROPRIETARY DATA

Telemetry & traces

Operator behavior, task outcomes, and ROS execution tied to real hydraulic systems.

02 · MODELING

1,200+ GPU hours

Near-deterministic 16-DoF behavior captured as reusable execution intelligence.

03 · SIM FIDELITY

Under 10 cm (Isaac Sim)

Credible path from simulation to reliable task execution and policy validation.

04 · DISTILLATION

Every machine feeds back

Field data reinforces skill distillation into the next generation of policies.

End state: open integration below, proprietary machine intelligence above - the Lattice of robotics for heavy machines.

§ MACHINE MOAT · END OF DOCUMENT