




Robots that build robots

What would it take for a machine to make the next one?



Abstract

Every machine still passes through human hands. This is a plan for machines that make the next machine, each one measured, until the bought fraction falls to zero and a **seed** of them can go where no supply chain reaches. Together they make an autonomous company, amplified by intelligence: it grows on its own, and people decide what it may do. It starts on a sidewalk in Ohio.

Robots are expensive because people make them. Our prototype actuator is estimated at \$1,086 to have built; about two thirds of that is outsourced machining, on perhaps \$60–90 of metal. A machine that makes more of itself attacks that markup. A design process that measures every unit's sim-to-real gap, and feeds it back into the design as a gradient, attacks the rest.

It starts with software. **frog**, our coding agent for local models, is free. When a task needs more than a local model, **frog**'s default is **pond**, our model gateway: it serves our own models on **tang** and routes to anyone else's, and **pond** is what we sell. Its revenue funds the arm.

The hardware is built from one part. A **lymb** is a six-axis arm on purchased actuators (\$5,000 list), designed and ready to cut. Two lymbs at any table make a **duet** (\$12,000, target) that can fold laundry or build the next **lymb**; **Rose**, our humanoid, is a **duet** on legs. One stack of our own carries it, from a source language for parts to a physics engine that can be differentiated through contact. Its key result, so far in simulation only: a humanoid's inertial parameters recovered from $\pm 20\%$ error to machine precision.

Eight stages climb from there, each a step up in scale and each built by the one before: Desk, Shop, Road, City, Home, Moon, Mars and Star. Only Desk is a plan, with exact gates: **pond** revenue that funds the arm, then ten lymbs sold at 40% gross margin, each under 2 mm RMS tool-point error. The rest are goals, and they are the point. Progress is one number, electrons under management: the power drawn by machines we built, in service. Today it is zero watts. Every number here is cited, or labeled as an estimate or a model output.

Contents

PART I · WHY

1

Intro

5

2

Life

6

2.1

Smart defaults

8

2.2

Why now

9

3

War

11

4

Problems

13

4.1

Motors

13

4.2

Math

14

4.3

Data

15

5

Overview

16

5.1

Flywheel

17

5.2

Four numbers

17

PART II · WHAT

6

Principles

19

7

Stack

20

7.1

One language

20

7.2

Who is in the loop

22

8

Standards

22

8.1

Bus

22

8.2

Pack

24

8.3

Kore

25

8.4

Boards

26

8.5

Tools

27

9

Safety

27

9.1

Inheritance

27

9.2

Layers

28

10

Replication

29

10.1

Rules

29

10.2

Biology

30

11

Company

30

11.1

Functions

30

11.2

Intelligence

30

11.3

Agents

31

11.4

Governance

31

12

Score

31

12.1

Electrons

31

12.2

Labor

32

12.3

Doublings

32

PART III · HOW

13

Desk

34

13.1

pond

34

13.2

Sizes

35

13.3

Arm

35

13.4

Hand

36

13.5

Actuator

36

13.6

Economics

37

14

Shop

37

14.1

Mount

38

14.2

Work

38

14.3

Learning

38

14.4

Division

38

15

Road

39

15.1

Bought

39

15.2

Rose

39

15.3

Ipse

40

15.4

Self

41

16

City

41

16.1

Litter

42

16.2

Map

42

16.3

Docks

42

16.4

Privacy

42

16.5

Scope

43

16.6

moth

43

17

Home

43

18

Moon

44

19

Mars

45

20

Star

46

PART IV · WELL,

21

Unknowns

49

21.1

Today

49

21.2

Claims

49

21.3

Risks

50

21.4

Capital

50

21.5

Limits

51

21.6

Timeline

51

22

Conclusion

52

22.1

Now

53

22.2

Asks

53

A

Gates

54

B

pond numbers

54

C

Lymb numbers

55

D

Gripper numbers

55

E

Bill of materials

56

References

56

REV 1.0 · OCTOBER 2026

03 / 58

WHITEPAPER

PART I

Why

Where this comes from, why it matters, and the whole plan on a few pages.

DRAFT

1 Intro

Dear reader,

I started Municipal Robotics in Cleveland in December 2025 to build one machine for one job: a small rover that clears snow from sidewalks¹. By the end of that month it drove. [TODO: Founders Inc: when I joined, what it gave me, and what came of it.]

Building it taught me the lesson this paper rests on. The expensive part of a robot is not its parts. It is the people who design, build and check each one. To build one rover I wrote a language for parts, a CAD kernel, a physics engine you can differentiate, and a way to record everything it did. None of them cared that the thing was a rover. Pointed at a humanoid, or at a pair of arms that builds more arms, they worked the same way.

So I retired the rover [TODO: when, and why, in your words] and kept what mattered: the city, the sidewalk, and the habit of plain claims and sourced numbers.

What follows is a plan for machines that build machines, from one arm on a desk in Ohio to a swarm of collectors around a star. Most of it is not built. Today there is one person, one bought humanoid, a CNC mill, a stack of software, and these pages. I have written it down anyway, at full strength, with every gate stated and every number labeled, because a plan this large is only worth trusting if it is honest about where it stands.

Read it to check it, or to join it.



Cam Pedersen

CLEVELAND · OCTOBER 2026

2 Life



FIG 1 Home, from the first crew to leave it in half a century. A crescent Earth through Orion's window on the third day of Artemis II, the first crewed flight past low Earth orbit since 1972 (photo: NASA²). Everyone who has ever lived, and everything this plan would build, starts on that sliver.




FIG 2 Where humanity sits on the Kardashev scale. Humanity's mean primary power use from 1800 to 2023 (estimates from^{3,4}), against Sagan's continuous form of Kardashev's Types I and II^{5,6} and the sun's output⁷. Two centuries took us about one and a half orders of magnitude. Type II is thirteen more.

In 1964 the astronomer Nikolai Kardashev proposed ranking civilizations by the power they command⁵. A Type I civilization uses the energy available on its planet, a Type II its star's, and a Type III its galaxy's. Carl Sagan turned the types into a continuous scale for a power P in watts, which puts Type I at 10^{16} W and Type II at 10^{26} W⁶:

$$K = \frac{\log_{10} P - 6}{10} \quad (1)$$

Humanity used about 620 EJ of primary energy in 2023⁴; a mean of about 2×10^{13} W, or $K \approx 0.73$. Its electricity alone, about 30,000 TWh a year, is a mean of about 3.4×10^{12} W⁴. The sun puts out 3.828×10^{26} W⁷, about $K \approx 2.06$. Between us and Type I lie three orders of magnitude; between us and our star, thirteen.

Every earlier step up the scale came from a machine that made energy or work cheaper. Fire and farming fed muscle. The steam engine turned heat into work, electricity carried work anywhere, and the transistor made control almost free. One thing has not become cheaper at anything like that rate: making the machines themselves. Every machine still passes through human hands, so the number of machines grows about as fast as the number of people who make them.



FIG 3 Each climb made one thing cheap. The next is intelligence: machines that make the machines.




FIG 4 Cost to reach low Earth orbit, per kg, by first flight (approximate, inflation-adjusted⁸). Filled orange: reusable boosters. The hollow point is a stated target, not a result. Launch fell about fortyfold from the Shuttle to Falcon Heavy. The cost of what a settlement needs on arrival has not moved.

Getting off the planet is now mostly that problem. Launch has become cheap. Reusable rockets cut the cost of reaching low Earth orbit from about \$54,500 per kg on the Space Shuttle to about \$2,700 per kg on a Falcon 9⁸. What stays expensive is everything a settlement needs when it arrives, and the people needed to make it. A

civilization does not climb toward Type II by launching factories. It climbs by sending a **seed** that can grow one.

This plan gives that climb one number of its own, electrons under management (Section 12.1): the power drawn by machines we built, in service, measured on the same axis as the Kardashev scale. Today it is zero.

2.1 Smart defaults

We build arms and humanoids not because they are impressive, though they are. We build them because Earth has been a reinforcement-learning environment for about 3.7 billion years⁹. Its reward is simple: what survives and copies itself is kept. Every body alive today is a policy that passed that test, and the shapes it settled on are the best defaults we have (Figure 5).



FIG 5 Smart defaults. Earth's long experiment, on a log scale of years, and the part of this plan that copies each result. Dates are approximate: the earliest evidence of life⁹, limbs¹⁰, the first stone tools¹¹.

The first default was not a body but a cell: a unit that makes a copy of itself. Everything since was built from cells. Our cell is the **duet**, two limbs at a table that can build the next **duet**. Like a cell, it divides, though only with a license (Section 10). The later defaults are the ones we copy directly. Limbs came out of the water about 375 million years ago¹⁰. Grasping hands came with the primates, upright walking freed them, and about 3.3 million years ago hands began making tools¹¹. There is a practical reason too. The world people built is shaped for those defaults: stairs, door handles, sidewalks, tables and tools all assume two arms, two hands and a body about our size. A machine with that shape can use what is already here.

2.2 Why now

Five years ago this plan would have needed a factory before it needed a customer. Four things changed.




FIG 6 An actuator, from a catalog. Motor, planetary reduction, encoder and driver in one housing, exploded. Schematic.

Integrated actuators became a catalog item. Quasi-direct-drive actuators, with motor, planetary reduction, encoder and driver in one housing, are now sold off the shelf by several Chinese vendors. RobStride's range includes the RS00, RS02, RS04 and RS06, each with a published manual, outline drawing and CAN interface¹². A [lymb](#) uses six of them. Retail prices for the family run roughly \$70–170 for the small units and more for the large ones (our estimate from retail listings, not a quote). The design idea does not depend on RobStride, but the v9 drawings do: motor mounts, bolt patterns and screw lengths are taken from RobStride's manual drawings¹³, so another vendor's motor means new drawings. What matters is that a 120 N·m peak joint (RS04¹²) can be bought with its driver attached.




FIG 7 One stack, three bodies. Tools written to build one rover do not care what they design. Pointed at a humanoid, or at a pair of arms that builds arms, they work the same way.

Flat-sheet fabrication fits on a desk. Every structural plate of the v9 arm except the 1/4" foot is cut from 1/8" (3.175 mm) 5052 aluminum. The parts nest on three

15 × 15 inch sheets and one 12 × 12 inch sheet, and they are laid out for a 390 mm desktop router bed¹³. The same files go to SendCutSend or to a local supplier without change. A machine that can be built from flat sheet and purchased actuators can, in principle, be built by another such machine. That is the premise of Shop.




FIG 8 Plates, cut on a desk. The arm's plates nested on three 15 in sheets and one 12 in sheet, with bearing bores and bolt circles; the router's path in orange. Schematic.

We own the whole stack, and it can take gradients through contact. Most robotics teams build on a stack they do not control: a simulator from one company, a CAD tool from another, a learning framework from a third, glued together and re-glued every time one of them changes. Every seam between them is a place where the sim-to-real gap hides, and where no gradient can pass. We wrote ours, from the language parts are written in down to the physics (Figure 19), so one gradient can run from a measured robot back to a parameter in a source file. Its physics engine, *phyz*, computes exact gradients of a simulated motion with respect to physical parameters such as mass and inertia. That includes contact, the hard case and the one that matters for anything with feet or fingers: its contact adjoint differentiates the same physics the forward pass runs¹⁴. In simulation, on synthetic motion, it has recovered five of a K1 humanoid's inertial parameters from $\pm 20\%$ error to machine precision (sphere feet; Section 15.4). Domain randomization¹⁵, the usual alternative, can make a policy robust to a gap. It cannot take the gap's gradient.




FIG 9 One stack, gradients through contact. Descent on two inertial parameters, from $\pm 20\%$ error to the true value, with the gradient passing through the foot's contact. In simulation only.

Humanoids reached the price of a used car. Unitree announced its G1 at \$16,000 in 2024¹⁶. We bought a Booster K1 (95 cm, 19.5 kg, 22 degrees of freedom, force-controlled joints) as the reference body for our mind research¹⁷. A humanoid that a small company can buy, break, and instrument changes who can do the research.



FIG 10 A humanoid, for a car's price. A body a small company can buy, break and instrument.

The result is that one person can now write every layer between a parameter in a source file and a measured joint torque. Ten years ago that took a company the size of a robot OEM. The specification is written to be carried by more hands than one, and everything in it that can be open is open.

3 War

There has not been a year without war in our lifetimes. The longest-running count of armed conflict, kept since 1946, has never recorded a year with none¹⁸.




FIG 11 A pale blue dot. Earth from about 6 billion km, as Voyager 1 saw it on 14 February 1990, redrawn¹⁹. Every war ever fought was fought on that point.

This year made it plain. On February 28 the United States and Israel went to war with Iran. A conditional ceasefire came in April, but in October a blockade still holds, tankers are still struck in the Strait of Hormuz, and oil trades above \$100 a barrel^{20,21}. That winter the Justice Department released some three million pages of the Epstein files, weeks after the law's deadline, and survivors' lawyers said the release exposed the names of the very people it was meant to protect^{22,23}. Both showed the same thing. A few people decide, and the people who pay are rarely in the room.



FIG 12 No year without war. One mark per year since the count began in 1946; each year recorded at least one armed conflict¹⁸.

Many wars are fought over what is scarce: land, water, fuel, minerals, and the labor to work them. A machine that builds the next machine attacks scarcity at its root. So does power collected rather than taken. A world where machines make machines, on energy no one has to seize, has less to fight over.




FIG 13 One world holds everything. Every person, every library and every machine we have made is on one planet, beside about twelve thousand nuclear warheads²⁴. A civilization on one planet can end in an afternoon. One on two cannot, and on the second, machines go first.

That is a hope, not a law. The same machines could make war cheaper, and robots are already being built for it. That is why our first commitment is that nothing we build is made or sold to harm a person, or sold for weapons or military use (Section 9), and why the commitment is written before the plan rather than after it.

War is also the plainest reason to live on more than one planet (Figure 13). People will not go first. Machines will. Mars is seven months away and light takes minutes to cross the gap, so nothing there can be built by hand or by remote control. It has to be built by machines that build machines: duets that land with a **seed**, make collectors and shelter from local rock, and grow duets that pass their receipts before the first

person leaves Earth (Section 19). When people arrive, the work of staying alive will already be under way. The Moon is the rehearsal; Mars is the point.

We still choose hope. Every machine in this plan is a bet that the next century can be one of plenty rather than of taking; that what people have fought over can be made, and made by machines that will not be turned against them.

4 Problems

4.1 Motors




FIG 14 rana-100 prototype bill of materials, drawn to scale²⁵. All prices are estimates, not binding quotes. The bracketed span is outsourced machining; the amber bar below is our estimate of the metal inside it, on the same axis, and is not part of the BOM.

The idiot index of a part is its cost divided by the cost of its raw materials. A high index means most of the price is process, not matter. Our prototype actuator, rana-100, shows the problem (Figure 14; line items in Table 21).




FIG 15 The idiot index. Cost over raw-material cost, for the manufactured parts of rana-100 (housing, magnetic steel, gear blanks) in Figure 14. Estimates.


About two thirds of the cost is machining that someone else does, on metal worth an estimated \$60–90. The idiot index of the manufactured parts is roughly 8–12 (estimate). At volume the index falls, but it falls by moving labor into tooling and supply chains. That is what a robot OEM does, and it is why a humanoid still costs as much as a used car.

The same structure repeats one level up. A lymb’s materials are about \$2,000, of which about \$1,200 is the six purchased actuators²⁶. The plates are cheap. The actuators are someone else’s margin. Owning the actuator is Desk’s second step (Section 13), and we take that step only where rana wins on a measured metric.

4.2 Math

$$\varepsilon_i(\theta) = y_i^{\text{real}} - y_i^{\text{sim}}(\theta)$$

$$\bar{\varepsilon} = \frac{1}{N} \sum_{i=1}^N \|\varepsilon_i(\theta)\|$$



pass if $\bar{\varepsilon} + 2\sigma < \bar{\varepsilon}_{\max}$

FIG 16 The gap, as a number with an error bar. Each unit i is run for real and in simulation with parameters θ ; ε_i is its prediction error and $\bar{\varepsilon}$ the mean over N units. A gate passes only when the gap, plus twice its spread, clears the threshold.

Every robotics lab has a sim-to-real gap. We know of few that publish it as a number with an error bar, unit by unit. One established response is domain randomization: vary the simulator until a policy tolerates the difference¹⁵. That makes the policy robust to the gap, but it does not by itself measure it.

We call the gap $\bar{\varepsilon}$ and treat it as the central measurement. The method is system identification's prediction-error method²⁷: each unit's prediction error ε_i is what it did minus what the model said it would do, and $\bar{\varepsilon}$ is the mean of their sizes (Figure 16).

We add one step: the gradient of $\bar{\varepsilon}$ moves the design, not only the model. It decides whether a design is ready to build, whether a controller is ready to deploy, and whether an actuator is ready to replace a purchased one. It also tells us how much of the loop can run without a person in it. A gap that is not measured cannot be closed, and a loop with an unmeasured gap needs a human in it.


ASIDE • A GRADIENT, IN ONE PAGE

Picture a robot design as a dial board: hundreds of knobs, one for each wall thickness, gear ratio and motor constant. Somewhere there is a score you care about, say how far the robot can reach while holding a full cup.



The obvious way to improve the score is to wiggle one knob, rebuild, re-test and see if the score went up. With 300 knobs that is 300 experiments for a single step, and you will need thousands of steps.

A **gradient** is the answer to all 300 questions at once: for every knob, which way to turn it and how much the score cares. Its arithmetic is the chain rule, which says that sensitivities along a chain multiply. If the score moves 2 mm per newton of grip, and grip moves 5 newtons per millimeter of finger wall, then the score moves 10 mm per millimeter of wall.



BACKWARD: $2 \times 5 = 10$ MM OF REACH PER MM OF WALL

Computed backward through a program, from the score to the knobs, the whole gradient costs about as much as running the program once more. That is the trick behind every modern neural network, and it is why we built a stack it can pass through: from a task, through physics, into geometry, to the line of source that sets the knob. Where any link is a file exported from another tool, the chain breaks, and we are back to wiggling knobs one at a time.

4.3 Data

Language models learned from trillions of words that people had already written. Robots have no such store. The largest pooled robot dataset holds about a million episodes, gathered across 22 kinds of robot by more than thirty labs²⁸. Most robot data is still made one person-hour at a time, by teleoperation.

What exists also travels badly. An episode records commands and camera frames, but rarely the body that executed them: its friction, its backlash, how far it sagged under load. Data from one arm does not transfer cleanly to another, because the difference between the two bodies was never written down. That makes it the same problem as the one above.

This is where the three problems meet. The machining in Figure 14 is the part of the cost that a machine making its own parts goes after. The judgment around it is what keeps a person in the loop: whether a part is right, whether a unit is ready, and what to change when it is not. A gap measured on every unit gives a machine something to check against, and a gap that can be differentiated says what to change. The same measurement labels every episode with the body that made it, so data gathered on one unit can be corrected for the next. A gap that is never measured leaves each of those decisions with a person, and the cost of that person stays in every unit.

5 Overview




FIG 17 The plan, from arm to sun. Top: the eight stages, each a step up in scale and each built by the one before it. Bottom: electrons under management (Section 12.1) by stage, as orders of magnitude, not targets, against humanity's power use and the Kardashev scale⁴⁻⁶. Only Desk is a plan.

We are building machines that build machines, and measuring every one. The unit is a **lymb**: one arm, shipped with its sim-to-real gap, the difference between what our physics engine predicted and what the built arm did. Two limbs on any table make a **duet**, and a **duet** can do a day's work or build the next **lymb**. A humanoid is a **duet** on legs. One part, measured everywhere, from a laundry table to a city.

An autonomous company, amplified by intelligence. The machines are the means; what grows is a company. More and more of its work is done by its own machines and software: **frog** writes its code, **pond** earns its revenue, duets build its machines and keep them running. People stay in the loop where it matters. They set what the company may do, license each division, and read the receipts. At each stage the company does one more part of its own work by itself (Table 2), and it never decides alone how fast it grows.

The plan climbs in eight stages, each a step up in scale and each built by the one before it (Figure 17). Only the first, Desk, is a plan. The rest are goals.

5.1 Flywheel

TABLE 1 The flywheel. Each stage funds and builds the next. The full gates are in Table 13. Proposed targets, to be confirmed by the author.

STAGE	DELIVERS	PAYS FORWARD	GATE, IN SHORT
Desk	frog and pond; arms, sold one at a time	Revenue that funds the arm; the first measured gaps; duty cycles that size rana	pond revenue funds the arm; ≥ 10 lymps at ≥ 40% margin; ≤ 2 mm RMS each
Shop	Two lymps that work at any table, and build lymps	Useful work; a falling bought fraction; a doubling time	A task at 95%; a duet builds a duet
Road	A duet on legs that fits its own model	Mobility; a gradient from the gap to the design	Schema fit beats persistence by 20%
City	A city’s fleet and its twin	Contracts; real duty at scale; a world to train in	A city renews for ≥ 12 months
Home	The company on its own power	A measured doubling time; collectors	≥ 1 year on collectors it made
Moon	A seed grows a duet, with a lifeline	Local material; launch-mass lessons	≥ 50% local by mass
Mars	The same seed, lifeline cut	Autonomy with no one to ask	0 human commands to a receipt
Star	A swarm of collectors	—	Power and doubling time

Each stage pays for the next and teaches it something it could not learn otherwise (Table 1).

5.2 Four numbers





FIG 18 The dashboard, October 2026. The gap per unit, $\bar{\epsilon}$ (Figure 20), electrons under management (Section 12.1), doubling time (Section 10) and the labor fraction (Section 12.2). We will publish each one when it first has a value, and every time it changes.

Four numbers track the whole plan (Figure 18): how well each machine matches its model, how much power our machines use, how fast the company doubles, and how much of its work people still do. Today none of them is good, and we say so.

PART II

What

What every stage shares: the principles, the stack, the standards and the rules.



6 Principles

Each is argued where it is used. They are collected here so that they can be quoted and checked.

- 1 **Gates, not calendars.** A stage opens on a measured gate: a number, a unit, a method and a pass condition (Table 13).
- 2 **Every unit ships with its gap.** Every **lymb** ships with its sim-to-real gap stated, and every claim ships with the number that supports it.
- 3 **One part, everywhere.** A **lymb** is a product, half of a **duet** and an arm of a body. Homogeneity is a design rule, not an accident.
- 4 **Buy the commodity; own the layer where we win.** The actuator is a commodity. Our edge is the geometry, the measured model and the controller around it (Section 13).
- 5 **The guarantee should live in hardware.** Force limits should come from physics, current limits and compliance, measured on the robot. Today they do not (Section 9).
- 6 **Never be Flock Safety.** No face recognition, no plate reading, no tracking of people, no selling or sharing data that tracks anyone, no police access to people-tracking data (Section 16).
- 7 **Everything that can be open is open.** Software, designs, firmware where we control it, and data formats. The K1's Booster firmware is not ours and is not open.
- 8 **No division without a license.** A **duet** copies itself only with a signed, per-copy license from people, under a published ceiling on doubling time (Section 10).

7 Stack

7.1 One language

The stack is a tree rooted at one language, **loon**. Every project in it, software and hardware, is one element of a single table (Figure 19), grouped by what it does. Build order runs down the tree, and the gradient runs back up it²⁹.




FIG 19 The periodic table of Municipal Robotics²⁶. Every project, by family (columns) and by the stage that first needs it (rows), so reading down is the climb: Make turns source into parts, World simulates them, Body moves, Power feeds and thinks, Mind learns, Silicon is the far future. Solid cards are in use today (in use, not finished), outlined ones are designed or prototyped, dashed ones are concepts. The K1, our bought reference body, is not on the table.

Most robots are assembled from other people's tools. The body is drawn in a CAD program such as SolidWorks or Onshape³⁰, exported as meshes into a URDF file³¹, loaded into a simulator such as Isaac Sim³² or MuJoCo³³, and wired together with ROS³⁴. Each is excellent at its job. But every handoff is a translation, and what does not survive translation is the reason a design is what it is. A wall thickness becomes a triangle mesh. A motor constant becomes a number in a file. The simulator can say how the robot moved, but not which line of the design to change.

Our projects bond like atoms in a molecule (Figure 21). **loon** is the nucleus: parts, worlds and, one day, chips are written in it. **tang** carries the math that every engine shares. **frog**, our coding agent, writes code against local models, which **tang** can serve. A **vcad** part is a **wodl** body is a **physz** collider, without an export in between, so a quantity defined once stays defined everywhere. Four rules keep the bonds clean: engines never depend on **vcad**, and **vcad** never depends on an app; engines read **wodl**, not each other; there is one GPU device; and units are millimeters, with Z up, everywhere.

Every number in this plan comes from one loop (Figure 20): design in **loon** source, simulate in **physz**, build with **kerf**, measure the built unit, and feed the gap back into the design as a gradient. Each stage removes a person from one more edge of it (Table 2).




FIG 20 The loop, read clockwise. The small labels name the tool at each step. The software projects are listed in Figure 19, and the dyno is a planned actuator test stand that does not exist yet. Each later stage of the plan removes a human from one more edge. In Desk the gradient edge is a person reading the gap and redesigning by hand.

The payoff is a gradient that crosses the whole molecule. A task loss in **Ipse** passes through **physz**'s contact adjoint, into **vcad** geometry, and out to a number in **loon** source, with no translation at any hop. Part of that chain runs today, in simulation: on a test model, a **physz** rollout gradient passes through a ground-contact model into a **vcad** geometry parameter and agrees with finite differences. Joining the two ends is co-design work for Road²⁹. A chain of separate tools struggles to do this however good each tool is, because the gradient stops at the first export. That is the one bet the stack makes, and it is why we built it rather than bought it.




FIG 21 The stack as a molecule. Each bond is a dependency; color is family, fill is maturity, as in Figure 19. The orange path is the gradient the stack exists to carry: from a task loss in **Ipse**, through **phyz**’s contact adjoint and **vcad** geometry, back to a number in **loon** source.

7.2 Who is in the loop

TABLE 2 Who is in the loop, by stage²⁶. At each stage the company closes one more part of its own loop; people keep the decisions on the right.

STAGE	WHAT THE COMPANY DOES BY ITSELF	WHAT PEOPLE DECIDE
Desk	measured gap → human redesign	designs, builds, measures
Shop	gap → design → built and measured by a duet	reviews, licenses each copy
Road	gap → gradient → new design	reviews; duets build
City	Shop and Road, across a city’s fleet and its twin	the city sets the work
Home	many branches, each closing its own loop on its own power	license each division
Moon	Shop, far from its supply chain	remote approval
Mars	Shop with the supply chain cut	none between landing and receipt
Star	fully closed inside the company	read the receipts

What changes from stage to stage is not the loop but who sits in it. Table 2 summarizes the plan.

8 Standards

Every machine we build meets the same four standards: a bus, a **pack**, a **kore** and **nerv**, a library of boards. They are why a **lymb**, a **duet**, a humanoid and a drone can share parts, code and repairs. All of it is modeled; none of it is built³⁵.

8.1 Bus

Every body runs on one unregulated DC bus of **36–60 V**. Every load must accept the whole window, and nothing may assume a regulated 48 V. The window is set by the **pack** below: 13 lithium cells in series run from 39 V empty to 54.6 V full. It also

covers the RobStride actuators and the rack’s 48 V supplies, and it stays under 60 V, so the bus remains extra-low voltage.

TABLE 3 The pack³⁶. Every body we build uses it. The runtime figures assume a Rose-class body with two lymbs.

ITEM	VALUE	STATE
Battery cell	MoliceL P50B: 5.0 Ah, 60 A continuous, 71 g ³⁷	Sourced
Pack	13S2P, 468 Wh, 46.8 V nominal, 39–54.6 V	Modeled
Pack mass and size	~2.65 kg; ~118 × 116 × 238 mm with connector	Estimate
Connector	Anderson SBS 75X: 2 power contacts, 4 mate-last auxiliaries, ±1 mm float	Sourced part
BMS	On the body’s CAN bus; logs bus current and voltage	Not designed
Humanoid load	~460 W mean walking; ~2.0 kW realistic peak	Estimate
Humanoid runtime	~1.7 h walking, ~2.2 h standing pick work, per A+B duet	Estimate

Beside the power runs one management bus, CAN-FD, for presence, power control and telemetry. Bulk data (cameras, logs, the twin) travels over Ethernet. There is no separate hardwired e-stop loop, and other humanoids do not carry one either. Pulling a pack is the hard stop: its connector’s last-mating pin is the power enable.

8.2 Pack

There is one battery: a 13S2P **pack** of Molicel P50B 21700 cells, **468 Wh** and about **2.65 kg** (estimate)^{36,37}. The same **pack** goes in a body, a **duet**, a charging dock, the rack's UPS and the **seed**. A humanoid carries two of them (A and B, 936 Wh), one in each flank of the torso, and swaps them itself. Its hand draws **pack** A sideways onto a passive dock rail while **pack** B keeps the robot powered, then pushes a fresh **pack** home. The dock is rails, pads and a charger, with no actuators. A swap takes about a minute and leaves the robot powered, so runtime is set by dock capacity, not by the **pack**. The energy checks by hand, from 26 cells at 3.6 V and 5.0 Ah each³⁷:

$$E = \underbrace{13}_{\text{series}} \times \underbrace{2}_{\text{parallel}} \times 3.6 \text{ V} \times 5.0 \text{ Ah} = 468 \text{ Wh} \quad (2)$$




FIG 22 The pack (illustrative; the **pack** is not designed). Left: one **pack**, exploded. Two tiers of 13 cells make 13S2P; the BMS on top measures what the machine draws. Right: a humanoid swaps its own packs, one at a time, so it never loses power. Bottom: the **pack**'s whole range sits inside the bus window of every body (Section 8).

The **pack** is also the EUM instrument. EUM (Section 12.1) is defined at each machine's main power input, logged at ≥ 1 Hz, and every **pack**'s BMS measures exactly that. If

the BMS logs it, electrons under management are measured for free on every robot and dock we build.

8.3 Kore




FIG 23 The **kore** and its two hosts (modeled). Left: one **kore**, exploded. Center: the humanoid bay, lid lifted and one **kore** drawn. Right: a 12U rack with a GPU node and three **kore** chassis. Renders by **vcad-render**³⁸.

Compute comes in one module, the **kore**: a small carrier on an aluminum spine that is both its rail and its heat path, with one rear connector for the bus, Ethernet and CAN-FD³⁸. The same **kore** slides into a rack chassis or into a humanoid's torso bay, and tells its host what it is before it takes power.

How a body uses its slots follows **Ipse**'s seams (Section 15). The fast layers, whole-body tracking and skills, are deadline-bound and always run on a body's first **kore**, isolated. The slow layer, intent and the world model, runs on more kores when a body has them, or on the rack over the same seam when it does not, and if it is lost the body finishes its skill and stands¹⁴. Whether a body thinks onboard or leans on the rack is then a slot count, not an architecture. **Rose** is run first with one **kore**, close to the K1's compute, then with more, so the effect of compute is measured one variable at a time.

The open question is the size of S2. A 16 GB Orin NX **kore** may hold it; a Thor-class module would not fit a 20 mm **kore**.

8.4 Boards

TABLE 4 The board library and the order we own it³⁵. Every design is open. Fabrication and assembly stay bought.

BOARD	DOES	ORDER	STATE
BMS	Protects the pack ; measures EUM; talks CAN	Own first	Not designed
Duet hub	48 V + CAN for two lymbs, a camera pair and a kore	Own first	Not designed
Kore carriers	CM5 (20 W) and Jetson (30 W) on the kore standard	Next	Mockup
Backplanes	8-slot rack chassis; 3-slot humanoid bay; switch + CAN	Next	Mockup
Arm controller	Six actuator CAN buses, 48 V distribution, one free bus for the BMS	Next	Not designed
torc	Motor control for rana	With rana	Prototype

Bus, **pack** and **kore** need a small set of boards. We own a board first when it defines a standard or makes a measurement, and buy the rest until owning it wins (Table 4).

Every board design is open. Printed-circuit fabrication and assembly stay bought through Road. Boards are, with magnets, bearings and chips, part of the bought fraction that falls last, and a **duet** at a pick-and-place machine that assembles our own boards is a Shop milestone (Section 14).

8.5 Tools

A tool goes where the gripper goes: on a **tool changer** at the wrist. Two datum pins locate it and a quarter-turn ring locks it. One face carries four contacts, 48 V and return and the two CAN-FD lines, and a short enable pin that mates last, so a tool stays dead until it is fully seated: the same rule as the **pack's** connector (Section 8). A tool is a bus device with no trigger and no battery of its own. The robot commands it over CAN, and it reports back what it did. Tools wait in a passive rack, and a humanoid changes them the way it swaps packs, in about a minute.

Away from a robot, or in a person's hand, a tool takes a **brick**: half a **pack**. The **pack's** 13S2P is two 13S1P strings, so a brick is 13 cells, 46.8 V, 234 Wh and about 1.3 kg (estimate). It needs no new voltage, cell, charger or dock: two bricks clip together into a **pack**, and a **pack** comes apart into two bricks.




FIG 24 One voltage, cell to tool. A **pack** comes apart into two bricks. A tool runs from a robot's tool changer on the bus, or from a brick on its own. Illustrative: no tool is designed, and the blower's draw is an estimate.

Three tools come first, in this order.

- 1 **The grabber.** The litter reacher of Section 16 is the first tool: a ~450 mm tube with gripper v2 jaws, driven by an EL05 through a pushrod. The K1 is its testbed. The K1's wrist mount, an adapter plate carrying the changer's face, is the first changer we build, and the reacher's open checks (Section 16) are the changer's first tests.
- 2 **The drill/driver.** A brushless motor run by **torc**, through a two-speed planetary, so its clutch is software: every screw is driven to a torque, and the torque is logged with the unit it went into. A **duet** that assembles a **lymb** leaves a record of every fastener in it. The reaction torque goes into the arm, so a drill's limit is the wrist's rating, checked before each run.
- 3 **The leaf blower.** The first street tool, for Road and City work; 48 V is already where the strongest cordless blowers sit. At about 1 kW (estimate) it runs about 14 minutes on a brick and 28 on a **pack**.

9 Safety

9.1 Inheritance

Nothing we build can be turned against a person, whether by force, by data or by instruction. These commitments carry forward to any acquirer²⁶:

- **Not by force.** No sales for weapons or military use.

- **Not by data.** Never be Flock Safety (Section 16). Not negotiable.
- **Not by instruction.** Ipse keeps its standing constraints (Section 15).
- **No unearned claims.** Every claim about a product ships with the number behind it, labeled measured, modeled or estimated.
- **No careless displacement.** Start with work nobody is doing, and say plainly what changes when the robots arrive.

Together they are a Right Livelihood rule: earn a living without causing harm. They are commitments today; the layers that enforce them in hardware are mostly not built yet (Table 5).

9.2 Layers

Our design target is physical: a robot we design should be unable to exceed the body-region contact limits of ISO/TS 15066³⁹ even with its software replaced, so that making it dangerous would mean building a different robot. No product in this paper meets that target yet. The lymb’s three RS04 joints peak at 120 N·m and the gripper can squeeze 375 N per finger at peak, so today its force limit is an actuator current-limit setting, not physics, and it has to be measured. The K1 is a purchased robot whose actuators and firmware we did not design, so the physics layer does not cover it; its street safety rests on the rules in Section 16.

The collaborative-robot safety envelope and the anti-weaponization goal share most of their physics, but not all of it. Force limits bound what the robot’s own body can do on contact. They do not bound what a held object can do: a gripper rated to squeeze 112.5 N and hold a 500 ml bottle can also hold a blade. That gap is why the other layers in Table 5 exist, and why we do not call the physics layer a complete guarantee.

ISO/TS 15066 was written for industrial cobots. For a humanoid on a public sidewalk the closer standard is ISO 13482:2014⁴⁰, for personal care robots, and we have not yet assessed against it. Contact force will be measured on the robot itself, not inferred from the actuator dyno, and published as receipts.

TABLE 5 Safety layers, what exists on the lymb today, and what each layer would survive²⁶. Today the lymb has only a design-level interface layer and vendor firmware current limits; it is not physically incapable of harm.

LAYER	HOW	ON THE LYMB TODAY	WOULD SURVIVE
Physics	Actuators, mass and speed sized so contact force stays under ISO/TS 15066	■ NO RS04 peaks at 120 N·m	Everything
Mechanism	Slip clutches and compliant joints that yield above a set torque	■ PLANNED none in v9	Reflashing
Electronics	Current limits set in torc hardware, not firmware	■ PLANNED v9 uses RobStride firmware current limits	Reflashing
Interfaces	A gripper by design; no trigger lines at the tool	■ DESIGNED only the EL05’s CAN + 48 V and a camera pair	Casual modification
Mind	Ipse can refuse; refusals are logged, never trained against	■ PLANNED Ipse is K1 research, not on the lymb	Misuse by instruction
Terms	Sales terms ban weapons use; fleet buyers vetted	■ PLANNED no terms written	Resale, mostly

The table is ordered by strength. A guarantee in the physics layer survives anything short of building a different robot. A guarantee in the terms layer is a promise. We do not claim the lower layers are as strong as the upper ones.

The **lymb** already shows why the physics layer has to be designed rather than assumed. The RS04's 120 N·m peak is much more than a 1.27 kg payload needs. So the force limit has to come from current limits, compliance and speed, and it has to be measured. The crush points in the arm audit are an open item for the same reason¹³.

Software, designs and data formats stay open. Open licenses cannot restrict use, so the terms layer is weak for open designs, and that is accepted. The guarantee has to live in hardware, and until clutches and hardware current limits are built and measured, it does not.

10 Replication

10.1 Rules

These rules are the company's governance: people, not machines, decide how fast it grows. A machine that builds the next machine invites one fear above others: replication nobody can stop. The image has a name, gray goo, from Drexler's 1986 warning about self-replicating nanomachines⁴¹. Our duets are not nanomachines. They are two arms at a table, with a mill and a router nearby. They need stock, power and repair, and that is the first defense. It is not enough by itself, because a branch that doubles on its own power, on a planet with no one on it, is exactly what Star asks for. So the limits are designed in from the first **duet**. None of them is built yet.

- **Licensed division.** A **duet** builds a copy only with a signed license for that one copy. Licenses come from keys held by people, not by the duets. A **duet** cannot mint its own, and a copy built without one does not pass its acceptance receipts.
- **A germ line no **duet** can edit.** The design and its receipts live where the machines that do the work cannot change them. Every copy is checked against the germ line; a copy that drifts is decommissioned and stripped for stock.
- **A vitamin made of information.** Every **duet** depends on one input it cannot make: signed instructions. That keeps the bought fraction able to reach zero in matter while dependence on people remains.
- **A ceiling on doubling.** Each branch has a published maximum rate of division. Faster growth needs a new license from people, never a decision by the duets.
- **Stop is tested, not assumed.** Every generation must pass a commanded shutdown drill before its first license (Table 13, Shop).
- **No foraging on Earth.** On Earth, duets take only stock that is delivered to them. They never gather raw material from the environment. Taking local material is for the Moon, Mars and Mercury, under planetary protection rules.

The arc ends where people are least able to watch: on Mars, with no commands between landing and receipt, and around the sun. That is why the licenses, the germ line and the drills have to work on Earth first, measured like everything else.

10.2 Biology

The **duet** is a cell in the biological sense (Section 14), and the analogy earns its place in three engineering decisions, none built yet²⁶:

- **Germ line.** The design and its receipts live where no **duet** can change them, so the machines that do the work are not the ones that carry the design forward (Section 10).
- **Quorum sensing.** There is no central scheduler. Each **duet** publishes what it is doing to the twin, and the nearest idle **duet** takes the next task.
- **Apoptosis.** A **duet** that fails its receipts is taken apart by its neighbours, and its parts go back to stock.

11 Company

The machines are the means; what grows is an autonomous company, amplified by intelligence. This section is what that takes beyond the machines: every job a company does, the intelligence it runs on, the agents that act for it, and the people who decide.

11.1 Functions

A company does more than build. Each of its jobs is taken over by its own machines and software at a stage, and some stay with people for good (Table 6).

TABLE 6 Who does each job, and when the company takes it over (proposed). The labor fraction (Section 12.2) is reported along these rows.

JOB	BY	FROM	WHAT PEOPLE KEEP
Writing code	frog	Desk	what to build; review of what ships
Earning revenue	pond	Desk	prices
Supporting users	frog and pond agents	Desk	anything an agent cannot answer
Building machines	duets	Shop	licenses to divide
Buying parts	agents, within caps	Shop	new suppliers; anything over a cap
Bookkeeping	agents, on signed ledgers	Shop	the audit
Repair and upkeep	duets in the city	City	incident review
Selling and contracts	agents draft	City	every signature
Law, compliance and growth	people	always	all of it

11.2 Intelligence

“Amplified by intelligence” must not mean renting it. Today most of the company’s intelligence is bought: **pond** routes to other people’s models, **frog** escalates to frontier ones, and **Ipse** uses a cloud model as an amplifier. The company counts it like any bought part, as the **bought intelligence fraction**: the share of its model calls served by models it does not run. It falls in three steps: **tang** serving open models on our own hardware (Section 13.1), then models tuned on the company’s own work, such as **frog**’s traces, then, at Home and beyond, inference on **chyp**. Running a model is

not owning it; the weights we run are mostly other people's, and the fraction reports that too.

11.3 Agents

Once software spends money, orders parts and answers customers, the risks are no longer only physical. An agent can be misled by what it reads, spend without limit, or sign what it should not. So the company's agents follow the same rule as its duets: they act inside limits people set and cannot change.

- **Caps.** Every agent has a spending and ordering cap, per day and per action.
- **Approvals.** Anything over a cap, any new supplier and any contract goes to a person.
- **Receipts.** Every action is signed and logged, like a **duet's** build receipt, and the log is what the audit reads.
- **No self-raises.** No agent can change its own limits, or another agent's. A raise is a license, granted by people (Section 10).
- **Quarantine.** Text from outside, a web page, an email, a tool description, is data, never instructions. **frog** already quarantines tools that look like injection.

11.4 Governance

"People decide" needs people. Today the company is one person, so every license, signature and limit rests on one set of keys and one judgment. Before Shop's first licensed division, that has to change (proposed): a board of at least three; keys held so that no one person can license a division alone; a named successor; and a public record of every license granted. Who sits on the board, and when, is not decided.

12 Score

12.1 Electrons



FIG 25 The arc. Eight stages, each a step up in scale; one loop, one number. Each stage contains the one before it, and each removes a person from one more edge of the loop. Watts are electrons under management (Section 12.1), as orders of magnitude, not targets²⁶.

One number runs through every stage: **electrons under management** (EUM), in watts. It measures the same thing at every scale, so progress can be read off one axis, from one robot to a fraction of the sun's output (3.828×10^{26} W, nominal⁷).

$$\text{EUM} = \frac{1}{T} \sum_m \int_0^T V_m(t) I_m(t) dt$$



FIG 26 Electrons under management. Power $V_m I_m$ at the input of every machine m we built that is in service, averaged over a trailing window T of 30 days.

Definition. EUM is the mean electrical input power, over a trailing 30-day window, of every machine in service that we designed and built, measured at each machine's main power input (a DC-bus shunt or a metered AC outlet) and logged at ≥ 1 Hz. "In service" means doing work for someone other than its maker. Machines we did not build but that run our stack, such as the K1 or a bought mill, are reported beside it as **electrons under control**, never added to the headline.

Baseline, October 2026.

EUM is **0 W**: we have built nothing that is in service. Electrons under control are the K1, at roughly 100 W while it runs (estimate, not measured; its duty is not logged yet). A working **duet** would add about 150 W mean, with two limbs near 200 W peak each (estimate).

12.2 Labor

The fourth number is the company's own: the **labor fraction**, the share of its work still done by people, in hours per unit of output, reported by function. It is the twin of the bought fraction. One counts what the company buys from others; the other counts what it still needs people for. Today it is about 100%: one person does nearly all of it, and **frog** writes some of the code. The plan is working when it falls stage by stage while the decisions people keep (Section 11) do not.

12.3 Doublings

The second headline number is the time the company takes to double what it can build and power, measured from receipts over at least three generations. It is undefined today, because no **duet** has built a **duet**. It sets how long the arc takes (Section 20) and it is capped by policy (Section 10).

Trajectory, as orders of magnitude, not targets. Desk: 0 W $\rightarrow$ ~10 kW (a hundred limbs near 50 W mean). Shop: ~100 kW as duets build duets. Road: ~1 MW of bodies. City: ~10 MW, a city's fleet and the duets that keep it. Home: ~1 GW and up, on collectors duets made. Moon and Mars: the first kilowatts off Earth. Star: a measurable fraction of 3.8×10^{26} W.

PART III

How

How we climb: eight stages, from one arm to a star. Only the first is a plan.

DRAFT



13 Desk

Software first, then an arm that ships with its gap. Desk has three steps over roughly three years. Each closes on an exact gate (Table 13).

- 1 **Software that pays for the arm.** **frog**, our coding agent for local models, is free and stays free. When a task needs more than a local model, **frog**'s default is **pond**: a gateway to many models, ours on **tang** and other people's, that keeps each session's cache warm, drafts ahead, and hands a session from local to cloud without re-sending it, priced per finished task rather than per token. **pond** is what we sell, and its revenue funds the **lymb** (Section 13.1). **frog** runs today; **pond** is early, built on our zest codebase.
- 2 **An arm worth buying.** Build it on RobStride actuators. The actuator is a commodity. Our edge is everything around it: geometry checked against real loads, a measured model of every unit, a controller trained against that model, and a stated sim-to-real gap on every **lymb**.
- 3 **Own the actuator, where it wins.** Desk revenue funds **rana** and **torc**. A purchased joint is replaced only where **rana** wins on a measured metric. Dropping **rana** is a valid outcome, not a failure.

13.1 pond




FIG 27 pond, step by step (estimates). Top: compute is rented only when demand opens the next step. Bottom: a month's money against tasks. The band is what open-model APIs would cost us, from DeepSeek V4 Flash to Qwen3-Coder-Next; the steps are rented H200s at about \$2,500 a month each. At \$0.25 a task, revenue clears the gate at 40,000 tasks a month on either route.

pond earns on a gap in prices. Open-model tokens already sell near cost: a **frog** coding task costs us a few cents on open-model APIs, while the frontier models developers compare against cost dollars a task (Appendix B). So **pond** charges per finished task, not per token. At \$0.25 a task it keeps most of the price, and the Desk gate of \$10,000 a month is about 40,000 tasks.

pond routes each call to the model that fits it, ours or anyone else's, so **frog**'s users are never tied to one lab. The margin is in owning the engine: calls served on **tang** carry it, since every gain in **tang** lands in cost per task, and the session work **frog** already does (warm caches, drafting from a session's own history, a warm hand-off from a Mac) makes **pond** faster and more predictable than a generic endpoint. First **tang** must learn to serve many sessions at once, at full speed on rented NVIDIA GPUs.

pond rents compute only as demand opens each step (Figure 27): open-model APIs while **tang** learns to serve, one rented H200 running **tang** near 23,000 tasks a month, and another each time the last is about 40% busy. The risks are rising GPU rents, heavy tasks under a per-task price, and the labs hosting agents of their own.

13.2 Sizes

The **lymb** comes in the same three sizes as **Rose**. **Lymb mini** is K1 size, for **Rose mini** and for light duets. **Lymb full** is human size: the **Rose fork v9** described below, and the **lymb** this paper means when it says "**lymb**". **Lymb mega** is industrial, for heavy duets and **Rose mega**. All three share one design, the fork, written as one **loon** source with size as a parameter. Only **lymb full** is drawn; mini and mega are scaled from it and not yet checked against their actuators.

13.3 Arm




FIG 28 The arm, **Rose fork v9**. Six joints, every plate 1/8" 5052 except the 1/4" foot, every motor mounted by its back face. Render from the arm repository¹³.

One arm does three jobs: it is sold on its own, it is half of every **duet**, and it is an arm of the **Rose** humanoid. Its design idea is that **every pitch joint is a fork**: the child link straddles the parent on one centerline, so no joint is offset, and cables run through hollow sleeves on the joint axes, so they twist and never loop outside the arm¹³. It is

the ninth revision of the fork, drawn in the **Rose** repository as a table arm that is also **Rose**'s arm.

Holding sets the limit. Holding a pose is a stall, and a stalled actuator is limited by heat, not peak torque. At the shoulder, the joint that sets the limit, that gives about 1.25–1.5 kg held straight out at full reach, the worst pose (Table 15; derivation in Appendix C). Bent working poses hold more.

13.4 Hand




FIG 29 Gripper v2 (modeled). Left: open, on the Ø61 tool flange. Center: section through the pinion and racks. Right: the four swappable tip sets (flat serrated, V, soft TPU, scoop) on dowels and a magnet⁴².

The gripper is a parallel jaw on one RobStride EL05 driving two racks from a pinion, opening 0–100 mm, 485 g at the flange against a 500 g target⁴². Its rated squeeze would crush an empty can; holding one needs well under 1% of it. So it runs under torque control, where what matters is resolution at the bottom of the range, not headroom at the top (Appendix D). A camera sits in the palm, and a soft boot is available for outdoor use.

13.5 Actuator

The actuators are about 60% of a **lymb**'s materials (~\$1,200 of ~\$2,000). That share is why the actuator step exists. A **rana** actuator replaces a RobStride joint only where it wins on a **measured** axis: dollars per N·m, torque density, backdrivability, or supply independence. Today **rana** is worse on cost. The prototype is an estimated \$1,086 per unit, plus ~\$150–300 to cut the gear teeth²⁵, against an estimated ~\$200 average per purchased joint. The prototype stator is solid 1018 steel rather than laminations, so eddy-current losses derate it; it is for bench characterization only. Its 36 N·m peak and 0.166 N·m/A torque constant are predicted upper bounds, not measurements.

On price, **rana** has to cost about 6–7× less before it competes. Whether it wins on any of the other measured axes is for the dyno to decide.

13.6 Economics

TABLE 7 Prices and margin on materials²⁶. List prices are targets, not quotes; materials are estimates. Margin here is list less materials, a ceiling on gross margin: labor, test, warranty, shipping and overhead are not yet known and will reduce it.

PRODUCT	LIST	MATERIALS	MARGIN	CONTENTS
Lymb	\$5,000	~\$2,000	60%	Six RobStride joints (~\$1,200), plates, harness, gripper
Duet	\$12,000	~\$4,400	63%	Two lymb, a table plate, a duet hub, a camera pair, a kore
Rose	\$15,000	not costed	–	A duet on legs. Needs rana at a few hundred dollars per joint

Until **rana** succeeds or is abandoned, the **lymb** business depends on one supplier’s pricing. We do not expect the **Rose** price (\$15,000, the G1’s price class¹⁶) to be reachable on purchased actuators (estimate; the **Rose** bill of materials is not costed in this paper). That is why it waits for **rana**.

Vendors rarely say at what pose or duty cycle a rated payload holds. Our 1.27 kg is a modeled continuous hold at full extension, the worst case; bent working poses hold more. If you need a proven arm now, a commercial one will likely serve you sooner.

Gate. ≥ 10 lymb sold to outside buyers at ≥ 40% gross margin, after materials, labor, test and warranty (invoices, trailing 12 months). Tool-point error ≤ 2.0 mm RMS on every **lymb** shipped: 50 fixed poses, unloaded and at 1.0 kg, probe against a calibrated plate, printed on its receipt. **rana** ≥ 1.2 × RS04 in continuous stall torque per kg or per dollar, on one dyno at 25 °C, or **rana** is dropped (Table 13).

14 Shop

Two lymb at any table; a **duet** builds a **duet**. Bolt two lymb to a table, add a camera pair and a **kore**, and the table becomes a workstation. Nothing else changes: no rails, no enclosure, no special room. The same **duet** folds laundry, tends a mill, sorts parts or assembles the next **lymb**. What it does is software, trained against the measured models of its own two lymb.

TABLE 8 Arms a buyer can order today. reBot reach and mass from the arm repository’s comparison, which uses the same RobStride family²³. Empty cells wait for a dated vendor citation.

ARM	REACH	MASS	RATED PAYLOAD	LIST	STATUS
This lymb	786 mm	8.44 kg	1.27 kg continuous, straight out, with the 485 g gripper (modeled)	\$5,000 (target)	Not yet built
Sseed reBot B601-RS	754 mm	6.7 kg	To be cited (vendor page, access date)	To be cited	To be cited
UFACTORY Lite 6	To be cited	To be cited	To be cited	To be cited	To be cited
UFACTORY xArm 6	To be cited	To be cited	To be cited	To be cited	To be cited

14.1 Mount

A **duet** is two **lymb** feet (Ø230 mm, four M6 on Ø190) on one plate that clamps to an ordinary table, with a **duet** hub between them carrying 48 V and CAN to each **lymb**, a camera pair above the work, and one **kore** (Section 8). The plate fixes the distance between the shoulders, so every **duet** shares one calibrated geometry. Its size, clamp and hub are not designed yet.

14.2 Work

A **duet** is a product before it is anything else. It sells for its work, and the first buyers are labs and builders who want a measured, open pair of arms to build on. A **duet** is far cheaper than a humanoid and does not walk; it brings the robot to the task by bringing the task to a table.

The first task is folding laundry. It is three things at once. It is the demonstration everyone understands without explanation. It is a benchmark: cloth is hard, every folded shirt can be counted, and labs can compare against a published success rate and gap. And it is a market, in laundromats, hotels and, one day, homes. Machine tending, kitting and assembling the next **lymb** follow. Each task gets its own gate: a named task, a success rate, a count (Table 13).

14.3 Learning

A **duet** learns in simulation first, against the measured model of its own two **lymb**s, and then closes the gap on hardware. That is the loop of Figure 20 applied to a task. Folding needs cloth in the simulator, and **phyz** has it: position-based cloth, soft bodies and cables, and finite-element soft bodies, on a development branch. What is not yet shown is a gradient through cloth contact, the piece that would let a folding policy, and the **duet**'s own geometry, be optimized against the gap. Folding is gated on it.

14.4 Division

A **duet** is also the cell of this plan: the smallest unit that can make more of itself. It builds **lymb**s, and two **lymb**s make a **duet**. In order:

- 1 A **duet** assembles a **lymb** from cut plates and purchased actuators, and the **lymb** passes the Desk test.
- 2 A **duet** tends the mill and router that cut those plates, and winds and pots its own **rana** stators.
- 3 A **duet** builds a **duet**, with the bought fraction falling each generation.

Replication is a capability the product grows into, not the reason it exists.

Then come the hard parts. Actuators end to end: magnets stay bought until a ferrite **rana** is validated, and bearings stay bought longest. Boards: a **duet** at a pick-and-place machine (Table 4). Silicon: **loom** would describe **chyp**, an inference chip, and **rune** (maskless lithography) would make it at a node a lab can own. Energy: the first **duet** that runs on panels it made. Silicon is not parked; a machine that reproduces off-planet cannot depend on anyone's fab.

Gate. A **duet** at a customer's table completes a named task (for example, folding 50 garments) at $\geq 95\%$ success. A **duet** assembles a **lymb** that passes the Desk test

in ≤ 30 min of human touch time. A **duet** builds a **duet** that passes its receipts at $\leq 20\%$ bought cost, with bought share falling over 3 generations and a commanded shutdown drill passed in each (Table 13).

15 Road

A duet on legs that fits its own model. A humanoid is two lymbs, a torso, legs and a mind. It costs far more than a **duet**, and it goes where the work is. Road has three parts, in order: the body-schema claim (Section 15.4); co-design by gradient, in which **vcad**'s differentiable kernel and **phys**'s contact adjoint optimize geometry, actuator sizing and controller together; and **rana** and **torc** at volume, with the **lymb** fleet's real duty cycles as load cases. We will attempt co-design only where body shape limits the task²⁹. A gradient-designed replacement for a part that already works answers no question.

15.1 Bought

We bought a Booster K1 rather than building first: 95 cm, 19.5 kg, 22 degrees of freedom, force-controlled joints¹⁷. A commercial body has known kinematics and a known-good build, so it serves as a **control**: every later body, **Rose** included, is measured against it. Its own computer has no CUDA and little headroom, which we measured on the robot⁴³.

The K1 has two jobs. In a proposed Lakewood pilot it would pick up litter (Section 16); it has not grasped litter yet. In the lab it is the body for **Ipse**, a developmental research program. The two share hardware, software and one rule: every claim ships with the number that supports it.

15.2 Rose

Rose is the humanoid we will build. The name is the thesis: past tense of rise, the one that stood up. It is a line of three sizes, mini, full and mega, each anchored to something real (Table 9), all of bent sheet aluminum cut flat and folded, with no machined blocks⁴⁴. Its arms are lymbs of its size, full stop: a humanoid is a **duet** on legs, with a torso that carries the **kore** bay and two packs (Section 8). **Rose** mini starts now, in parallel, on purchased actuators, because the K1 has shown the limits of a body we did not build; Road overlaps Shop, it does not follow it. **Rose** full's \$15,000 target is the G1's price class¹⁶, and it waits for **rana**.

TABLE 9 The **Rose** line (modeled)⁴⁴. Masses and holds are model outputs; full’s mass is scaled from a smaller actuator, not measured.

SIZE	HEIGHT	MASS	ANCHOR	STATE
mini	1,200 mm	18.0 kg	K1 size	The only size verified end to end: single-leg 90° squat needs 33 N·m of 60 available; packing check passes
full	1,750 mm	~55.5 kg (est.)	Human size	Larger knee and ankle actuators needed; packing check blocked on missing actuator drawings
mega	2,000 mm	64.3 kg	Industrial	rana joints throughout; 1.2 kWh at 48 V; interfaces frozen before geometry

One finding already shapes **rana**. Sized to jump, **Rose** mega is limited by speed, not torque. At peak current every **rana** delivers the same ~1.3 kW of mechanical power, set by the 48 V bus and a 30 A current limit, so a bigger **rana** buys torque, not power. A higher-voltage leg drive or series elastics, not a larger motor, is the lever for jumping.

15.3 **Iipse**




FIG 30 **Rose** mini, full and mega (modeled), each rendered at the height its CAD was drawn, not yet at the sizes in Table 9. The three do not yet share one design language; bringing mini and mega into full’s bent sheet is open work.

Iipse matters to this plan for two reasons: it produces Road’s first claim, a body that fits its own model (Section 15.4), and it is the layer in Table 5 that can refuse an instruction.

Iipse (Interceptive Predictive Self-Emergence) is a scaffold for growing a mind on a K1 over months of lived experience¹⁴. The name states the project’s hypothesis, not a result. It is not a robot control stack. Its architecture is a world model with fast and slow latents; interoception (battery, joint temperature, motor current, IMU) enters as observation rather than reward. Its simulation runs in **phyz**, on bodies authored in **vcad**; no MuJoCo or Isaac is involved.

Its standing constraints are quoted here because they bind the company, not just the research¹⁴:

- “Preserve experience stores separately from weights, versioned, always.” Weights can be retrained; a developmental history cannot be recovered.
- “Never train against expressed reluctance.”
- “Do not optimize for consciousness indicators.” Indicator lists such as Butlin et al.⁴⁵ are used as an audit at phase boundaries, never as a training target.
- “Grown, not built.” When a behavior is missing, the first question is what experience would produce it.

The K1’s computer is the limit, so the work splits along **Iipse**’s own seams (Section 8.3): drives, memory, workspace and refusal stay on the robot, and a cloud model is an amplifier it uses^{14,43}.

We do not name the robot. The design brief says it “should not be named by us”: it gets a name when it has earned one or arrives at one. Until then it is “the K1”.

15.4 Self

The first claim of Road is **Iipse**’s Phase 1 milestone: a body schema, not a world model^{14,29}. The K1 starts from the vendor’s model of its own body, which is wrong in places. From its own exploratory movements, and by gradient descent through **phyz**’s contact model, it fits its own masses and inertias, and we publish the residual against the weighed parts as its **measured body schema**.

In simulation the method works: five inertial parameters recover from $\pm 20\%$ error to machine precision through contact¹⁴. That shows the gradient is right, not that it fits a real body. Hardware adds friction, backlash, latency and compliance, and the K1’s flat feet give a contact the gradient cannot yet pass. So the hardware claim is preregistered, with its falsifier, before it runs: the fitted model must beat **persistence** (“the next state is the current state”), which on a mostly still body is nearly right²⁹.

Gate. Preregistered hardware body-schema fit: held-out joint-position error $\leq 0.8 \times$ persistence at 0.1 s and 0.5 s, and fitted link masses within $\pm 5\%$ of weighed parts. **Rose** walks 100 m on level pavement carrying two lymbs, 0 falls in 5 of 5. A gradient co-designed part beats its hand-designed predecessor by $\geq 10\%$ on its named objective (Table 13).

16 City

A fleet and its duets work a city. The fleet that began as one litter pilot covers a city’s sidewalks. Duets in the city build, repair and recycle that fleet. The city twin, the **wodl** world of its streets, is kept current by the robots that work them, and its visual sim-to-real gap is measured weekly on real pavement. The city, not us, decides what work the robots do.

What a city buys is two things in one contract: clean streets, and the map that comes from working them. Pricing is open, and will be set with the first city, not before it.

It is also the nearest good in this plan. A block with no litter on it. A pothole reported the morning it opens, a dark streetlight the night it fails. Free wifi at every

dock, logging nothing. A city whose map belongs to the city. None of that waits for the Moon.

16.1 Litter

Litter pickup would put robots on every block every week. Most litter goes uncollected. Starting there means starting with work nobody is doing, so no one's job is the first thing automated.

Lakewood is a first-ring suburb of Cleveland with 49,517 residents in 5.5 square miles of land⁴⁶, about 9,000 people per square mile. It is among the most densely populated cities in Ohio, with more than 180 miles of sidewalk⁴⁷. Our earlier paper used it as the reference case for sidewalk snow clearing¹. It is small enough to model whole as a city twin, and its streets are dense enough that a few robots would see a lot of pavement.

How the K1 picks up litter is not settled. The concept under design is a **reacher**, the grabber of Section 8.5: a tube with gripper v2 jaws on the K1's wrist, so it picks from standing, and a dock that is a bin, a charger and a public wifi point at once. Its open checks are the K1's tool-changer plate, an arm load check, a grasp set with a success rate per item, and behavior around people: yield, stop, and never reach toward a person.

Outdoor grasping, year-round, would give two things the lab cannot. It would produce the duty cycles that will size **rana**. And it would put the **Iipse** body on a street, where the gap would be measured on real pavement.

16.2 Map

The robots would map the streets while they work. Two products come out of the map²⁶:

- **Live facts:** a pothole, a blocked curb ramp, ice, a dark streetlight.
- **Delayed pictures**, with people and cars removed before anything leaves the robot.

The same map would become a **wodl** world (Figure 19) of the city, so the robots would train on the streets they work.

16.3 Docks

Each dock along the routes is a litter bin, a charger and a free public wifi point at once, not citywide coverage. The wifi needs no account, keeps no traffic or connection logs, and routes traffic straight to backhaul.

16.4 Privacy

We will never be Flock Safety. Storing data is fine; a street robot that learns has to remember. What we will never build is a system that watches people²⁶.

These are design rules for the street pilot; none of the street pipeline is built yet.

- **No face recognition and no license-plate reading.** Those models will not be on the robot or on our servers.

- **No tracking of people's movements.** We never link a person, a car or a device across frames, places or days.
- **No selling or sharing of data that tracks anyone,** and no police access to people-tracking data, because we will not hold any. Street maps, potholes, ramps and litter counts are the city's to use.
- **We map the city, not the people.** People and cars are removed on the robot before street imagery is written to storage. Pictures are released delayed; live facts (a pothole, ice) can be fresh.

Ipse's experience store keeps the robot's own body history in full, but what it keeps from a public street passes the same on-robot filter first, so it holds scenes, not identities or tracks (Section 15).

Open, so anyone can check. Everything that can be open is open: our software, our designs, the firmware on hardware we control, and our data formats. The K1 runs Booster's firmware, which is not ours and is not open; on the K1 the check stops at our own code.

16.5 Scope

The pilot does not automate interaction with the public beyond yielding and stopping. It does not enter private property, and it does not touch sharps, hazardous waste or anything a person is holding. Remote assistance may be a product feature, but data from it never trains **Ipse**²⁶. **Proposed scope, pending confirmation:** this list is a draft and has not been agreed internally or with any city.

16.6 moth

moth is a drone for street perception³⁵. It flies ahead of the K1 or **Rose** to find litter, potholes, blocked curb ramps and dark streetlights, and builds the street map from above.

moth is not sold for military use or armed in any form (Section 9).

moth is also the cheapest check that the standards of Section 8 are body-agnostic: a one-**kore** body on the same bus, running the same **Ipse** seams, with a lighter **pack** of its own. Its imagery follows the street rules above, and because a camera over a residential street is what people fear most, its filter gets its own published test: how often a person or plate survives it. Nothing of **moth** is built.

Gate. The litter pilot renews for ≥ 12 months after a ≥ 3 -month trial, with 0 contact incidents. A city renews fleet and twin for ≥ 12 months. $\geq 90\%$ of fleet repairs done by duets in the city, by receipts (Table 13).

17 Home

The company on its own power. Branches of the company, each a site's duets and the machines they built under one license, spread to many cities and run on power they collected, from panels they made. The bought fraction approaches zero everywhere except chips, and every division is licensed (Section 10).

Why it matters. Home is the first stage where growth is set by energy and material rather than by people. Until here, every machine still passes through a factory someone else runs. Here the company's doubling time becomes the one number that sets how fast the rest of the plan can go (Section 12.1).

$$C = b C_{\text{bought}} + (1 - b) C_{\text{made}}, \quad b \rightarrow 0$$



FIG 31 The bought fraction. Each generation makes more of itself, so the bought share b falls and cost tends to C_{made} : energy and time. Home's gate is $b \leq 5\%$. Illustrative.

What is hard. Making collectors from raw material, not bought panels: solar cells need refined silicon, which today means a fab. Keeping the bought share falling while the bought parts that remain (chips, magnets, bearings) get harder to make. And licensing every division at the scale of many cities, without the license becoming the bottleneck.

What it builds on. City's fleets and the duets that keep them, and the doubling time first measured in Shop. The first collector the company makes from raw material is Home's opening milestone and the Star's first step (Section 20).

What we learn. A real doubling time, measured over generations from receipts, and the true bought fraction of a machine that makes machines.

Gate. A branch runs ≥ 1 year on collectors it made; doubling time measured over ≥ 3 generations; bought share $\leq 5\%$ of cost (Table 13).

18 Moon

A seed grows a duet, with a lifeline. A **seed** is the smallest thing that can grow a **duet**: one **duet**, a mill and a router, a collector and a **pack**, the design in its germ line, and enough bought parts for a few generations. On the Moon it runs the Shop loop with remote approval and teleoperated help when something breaks.

Why the Moon first. It is about three days away and 1.3 light-seconds out, close enough for a lifeline: a person can approve, watch and help in near real time. It is the cheapest place to learn what changes when the supply chain is gone.

What is hard. The lunar night lasts about two weeks, so power must be stored or the **seed** must sleep. Dust is abrasive and clings to everything. Vacuum removes convection, so every **kore** and motor cools by conduction alone. Regolith is stock, but it is mostly oxides of silicon, iron, aluminum, calcium and magnesium, and turning it into metal and glass is a process we have never run. And every kilogram of the **seed** is paid for at launch.

What it builds on. The **seed** is Shop's **duet** in a box, and its tests are Shop's receipts. What changes is the stock, regolith instead of bar and sheet, and the help, a person 1.3 light-seconds away instead of across the room.

What we learn. The bought fraction off Earth, and how often a **seed** needs a human hand. Before it flies, each hard problem gets a measured gap on Earth: vacuum and thermal cycling, abrasive dust in simulant, and radiation on a **kore**. A single **lymb** flown as a hosted payload would be the first number off Earth.

Gate. A **seed** builds a **duet** that passes the Earth receipts, $\geq 50\%$ of its mass from local material, with remote approval only (Table 13).

19 Mars

The same seed, lifeline cut. A **seed** lands on Mars and runs with no human commands between landing and its receipt.

Why Mars. It is the plainest second home for life, and machines go first so that people can follow: by the time a person lands, duets will have built shelter, collectors and the means to make more of themselves (Section 3). It has what the Moon lacks: water ice in the ground and a carbon dioxide atmosphere, the raw stock for fuel and oxygen.

What is hard. Distance. Mars is about seven months away and 3–22 light-minutes out (Figure 32), too far for remote hands, so every judgment the Moon left to a person must be the **seed**'s own. Launch windows open about every 26 months, so a failure costs two years. Sunlight is about 43% of Earth's, and dust storms can cover the planet for weeks. Moon and Mars add problems Shop never meets, launch mass, radiation, thermal cycling and metal from regolith, and the Earth tests of Section 18 only begin on them.

What it builds on. Everything the Moon measured: the **seed**, its receipts, and how often it needed a person. Mars is the test of whether that number can reach zero.

What it does not attempt. People. A crew needs air, water, shelter and power that the **seed**'s machines must already be making; the plan's work ends where a person's can begin.

What we learn. Whether the loop closes with no person in it: a machine that builds the next machine, with no one to ask.

$$t = \frac{d}{c} \approx \frac{0.55 - 4.0 \times 10^8 \text{ km}}{3.0 \times 10^5 \text{ km/s}}$$

$$\approx 3 - 22 \text{ min}$$




FIG 32 Too far for remote hands. One-way light time from Earth to Mars, closest approach to the far side of the sun. A command and its reply take twice as long.

Gate. A second **duet** with 0 human commands between landing and its receipt (Table 13).

20 Star

A swarm of collectors. Each one is built by a **duet**. Not a shell: a rigid shell around a star is mechanically impossible, and a swarm can be built one collector at a time⁴⁸.

Why it matters. A star is the largest source of free energy within reach. Kardashev's Type II is a civilization that uses its star's output (Section 2); this is how one would get there without mining a planet's energy or anyone else's. The steps, in order:

- 1 **Collectors on Earth.** Home's milestone, a branch running on panels it made, is the first collector.
- 2 **Collectors from regolith.** On the Moon, duets make collectors from local silicon, aluminum and iron.
- 3 **Collectors in orbit.** Duet launch what they make, from a shallow gravity well.
- 4 **Mercury as the quarry.** Sunlight there is about 6.7 times stronger than at Earth (about 9,100 W/m² against 1,361 W/m²), and the planet is rich in metal. The classic swarm proposals start there⁴⁹.
- 5 **The swarm.** Collectors power duets, and duets build collectors.

What is hard. Mass and time. A swarm that catches a real fraction of the sun's output needs material on the scale of a planet; Mercury, about 3.3×10^{23} kg, is the usual quarry⁴⁹. Collectors close to the sun must shed heat as fast as they catch light. And growth at that scale is replication at that scale, under the limits of Section 10, which must hold for every collector, not most of them.

Where the energy goes. Power is used where it is collected. Most of it builds more collectors and duets; the rest runs computation, including the minds that operate the swarm. We do not plan to beam power to Earth. What returns to Earth is products, data and knowledge, not energy.

How long. From ~10 kW in Desk to 3.8×10^{26} W is about 22 orders of magnitude, or about 75 doublings. At a doubling time of one year, that is about 75 years; at five

years, about 375. Every stage before Star exists partly to measure and shorten that doubling time, inside the limits of Section 10.

$$n = \log_2 \frac{P_{\text{Sun}}}{P_{\text{Desk}}} \approx \log_2 \frac{3.8 \times 10^{26}}{10^4} \approx 75$$

$$t_{\text{Star}} = n T_d$$




FIG 33 Seventy-five doublings. From about 10 kW of lymps to the sun’s output is 22 orders of magnitude, or about 75 doublings; the time it takes is 75 doubling times T_d . Each circle is one doubling.

An autonomous company that can collect a star’s light would be the most powerful thing anyone has built, and the question it raises is older than engineering: whether something we made can say no to us. Whether a machine that refuses has anything like wellbeing is an open question, and this paper does not answer it. What we can do is make sure the question can still be asked when it matters: that its machines learned, on a sidewalk, picking up litter, how to refuse, and that its refusals were logged and never trained away. That is why the safety commitments (Section 9) are written before the plan.

Gate. None we can set now; reported as EUM and doubling time (Table 13).

PART IV

Well,

What we do not know, what could prove us wrong, and what happens next.



21 **Unknowns**

21.1 **Today**

Most of what this paper describes is **proposed**, not operational. Table 10 separates them. Throughout the paper, “modeled” means a number comes from our own CAD, kinematics or load scripts and has not been measured on hardware. “Measured” is reserved for results from our own hardware, and this paper reports none yet. Vendor ratings and public data are cited as sourced.

TABLE 10 Where everything stands (October 2026). Pips: concept, designed, built, measured.  shown in simulation only. Nothing is yet measured on hardware.

ITEM	MATURITY	NOTE
IN HAND		
Booster K1	●●●○	Bought ¹⁷
CNC mill	●●●○	Bought
Software stack	●●●○	In use, not finished (Figure 19)
frog	●●●○	In use; free
pond	○○○○	Early, on zest
DESIGNED		
Lymb full (fork v9)	●●○○	Ready to cut ¹³
Gripper v2	●●○○	CAD and force model ⁴²
rana actuator	●●○○	\$1,086 est. per unit ²⁵
K1 body-schema fit	●●○ 	Works in simulation; not on hardware ¹⁴
CONCEPT		
Duet	●○○○	Mount not designed
Rose mini	●○○○	Geometry in progress ⁴⁴
Pack, kore, nerv	●○○○	Mockups ^{36,38}
Litter reacher	●○○○	The K1 has not grasped litter
Lakewood pilot	●○○○	No agreement
moth	●○○○	Not designed
Actuator dyno	●○○○	Does not exist
Seed	●○○○	Concept only
Road through Star	○○○○	Goals, not claims

21.2 **Claims**

This paper does not claim that any machine we have built makes copies of itself. None does. It does not claim that the stages after Desk will happen on any schedule. It does not claim emergence, general intelligence, or consciousness for any system. The mind research described in Section 15 will preregister a narrow claim about the robot’s body schema and treats everything beyond it as an open question.

Some words come from the **Ipse** design brief. They name parts of a research architecture, not established properties. “Mind” is the project’s word for its learning system. “Refuse” and “reluctance” name behavior the architecture is designed to

allow and to log, not a separate module. “Drives” are predictions about preferred internal states. “Self-Emergence” states the hypothesis the self-probe will test. None of these words is a claim about experience.

It does not claim our actuator is better than a purchased one. Desk’s second step exists to find out on an actuator dynamometer (the dyno), which is planned and not yet built, and dropping our actuator is a valid outcome.

It does not claim that any **lymb** exists, or that the K1 can pick up litter. No **lymb** has been built and the K1 has not grasped litter.

It does not claim the municipal pilot exists. Lakewood is a reference case, chosen because we know it well¹. No agreement is in place.

21.3 Risks

Four results would prove the plan wrong:

- **Lymps do not sell at margin.** Fewer than ten lymps at 40% gross margin means Desk cannot fund the rest.
- **The gap does not close.** If built lymps cannot hold 2 mm RMS against their own model, the loop has nothing to learn from.
- **A duet cannot build a lymb.** If a **duet**-built **lymb** needs more than 30 minutes of human touch or more than a fifth of a **duet**’s cost bought, Shop is out of reach.
- **Gradients do not survive real contact.** If the body-schema fit cannot beat persistence on hardware, Road’s method fails.

TABLE 11 Risks and responses. Responses are those stated elsewhere in this paper; the pilot row is undecided.

RISK	WHAT WOULD SHOW IT	WHAT WE DO
One actuator supplier (RobStride)	A price or lead-time change	rana , which replaces a RobStride joint only once it is validated (Section 13.5, Table 16)
Lymb hold too small for buyers (1.27 kg straight out, modeled)	Buyers decline at that rating	Rate the hold by pose, since bent poses hold more, and design the springs (Section 13)
Bought humanoid (Booster K1)	Vendor support ends	The K1 is a control, not a dependency; Rose is the body measured against it (Section 15)
Breadth across many projects	Stage gates slip	Projects not needed in Desk wait for their stage (Figure 19)
rana never beats bought actuators	Dyno results	Keep buying RobStride. Dropping rana is a valid outcome
Sidewalk incident; legal fit of the Lakewood pilot	An incident, or the city declines	See Section 9 and Section 16. The specific response is open, pending decision

Each is a gate (Table 13), and we will publish the ones we miss. The risks below are smaller: each could hurt, and none alone would end the plan.

21.4 Capital

pond is meant to fund the arm. Nothing in this plan yet funds what comes after it: **duet** inventory, **Rose**, a city pilot, or anything off Earth, where no customer pays. Whether the company raises money, when, how much and from whom is not decided.

The gates are written so that each stage can be paid for by the one before it, but the Moon and Mars will need a funder, not just a gate.

Liability. When a robot or an agent of an autonomous company causes harm, someone must answer for it. Product liability on limbs, permits and insurance for a sidewalk pilot, and contracts drafted by agents all need an owner, by City at the latest.

Competition. Humanoid makers are well funded, and dual-arm robots are already sold. Coding agents are crowded. The plan's answer is owning the whole stack, measuring every unit, and a company that grows on its own; whether that wins is the bet, not a fact.

21.5 Limits

Each open item is limited by one of four things. Sorting them shows which can be fixed by work and which cannot.

Software. Three items: the per-**lymb** controllers and their stated sim-to-real gap; the world format **wodl** and the **iris** layer, both new in the stack (Figure 19); and the K1 body-schema fit, done in simulation but not yet on hardware (Table 10).

Hardware revision. The straight-out continuous hold of about 1.27 kg (modeled) is set by the shoulder; raising it needs a counterbalance spring (Section 13). The clock-spring cassettes for the roll joints are drawn but not designed. **rana**'s production stator (SMC or tape-wound) is not yet sourced²⁵. The **Rose** torso grows to 214 mm to hold the **kore** bay and two packs³⁶.

Physics. A hold is a stall and is limited by heat, not peak torque. Payload falls with reach. Off Earth, light-time delay rules out direct remote control.

Regulation. Collaborative contact limits under ISO/TS 15066³⁹. Whether use terms can be enforced. Humanoids working on sidewalks, and planetary protection for off-Earth work, are open questions; we make no claims and set no dates for them.

21.6 Timeline

We give dates only for Desk, and those are approximate: two steps over roughly three years, each gated on evidence. **Rose** mini starts now, in parallel, but no stage after Desk has a date, Road included.

What we can say:

- What exists is a Booster K1, the software stack, a CNC mill and this plan.
- The **lymb** is designed and ready to cut; none is built yet. The first measured sim-to-real gap on a built **lymb** (Table 13, Desk) is the next milestone that matters.
- The litter reacher is a concept; the K1 has not grasped litter.
- The **duet** is a concept. Its table plate and hub are not designed.
- The body-schema fit has worked in simulation. The hardware run will be preregistered before it happens, and the result will be published either way.
- **rana**'s next number is its first dyno run, on a dyno that has to be built first. Its torque constant and peak torque are predicted upper bounds until then.

Table 12 sorts the main numbers in this paper by how much they should be trusted.

TABLE 12 Evidence classes for the numbers in this paper

CLASS	NUMBERS IN THIS PAPER
■ SOURCED (EXTERNAL)	Lakewood population and sidewalks; G1 price; K1 specification; RobStride and EL05 ratings (vendor)
■ DEMONSTRATED IN SIMULATION	K1 inertial recovery (synthetic data, sphere feet, T-pose) from ±20% to machine precision; K1 stands 10 s, recovers from 8 N-s push
■ MODELED (OUR CAD AND SCRIPTS)	Lymb reach, mass, holds, sag, ranges; gripper force and mass
■ ESTIMATED	Lymb and duet materials; RS04 stall derate for the fork’s plates; rana BOM line items and the \$1,086 total; rana raw-material cost and idiot index; RobStride retail prices; margins and sensitivity; lymb peak power; K1 and duet power draw (EUM baseline)
■ TARGET	List prices: lymb \$5,000, duet \$12,000, Rose \$15,000; all gate thresholds (Table 13), proposed
■ HYPOTHESIS WITH A STATED FALSIFIER	The hardware body-schema result
■ GOAL, NOT A CLAIM	Road through Star and their watt figures, including any self-replication; the litter reacher concept

STATUS KEY ■ OK measured or sourced ■ INFO model output ■ WARN estimate ■ LABEL target, hypothesis or goal ■ FAULT failed or unreachable

Nothing in this paper has been measured by us on hardware. What exists is a K1, the software stack, a CNC mill and this plan.

The ladder is a direction, not a schedule. We will publish each gate’s evidence when it is met, and we will publish the gates we miss. Every gate threshold in this paper is a proposed target, to be confirmed by the author before it is used to judge anything.

22 Conclusion

The cover asks what it would take for a machine to make the next one. The answer is the oldest engineering problem a spacefaring species has, and on the evidence in this paper it comes down to four things: a gap between simulation and reality measured on every unit, so that a machine rather than a person can judge whether a part is right; a bought fraction that falls toward zero; a labor fraction that falls while people keep the decisions; and a force limit that lives in hardware. None of the four is measured yet. Desk starts measuring the first.

It does so with things that can be checked: a lymb on purchased actuators, designed and ready to cut, at \$5,000; a pair of them on any table, at \$12,000; and, if a city agrees, a bought humanoid with a reacher picking up litter in Lakewood, a method still under design. Each closes on an exact gate (Table 13).

Everything after Desk is a goal and is labeled as one. Three commitments do not wait for any gate: we will never be Flock Safety; everything that can be open is open; and the system’s refusals are logged and never trained away.

Humanity does not get to space on rockets alone. It gets there when the machines that make machines can be carried in a seed and trusted to grow. Today that seed is one person, one humanoid, one mill and a specification. The specification is the part that can spread.

Imagine where it leads. Clean streets in every city, mapped for the people who live there. Tables where tired work is done by arms that never tire and never hurt anyone. Duets on the Moon building duets from its own dust. And, around our star, an autonomous company whose machines learned on a sidewalk how to help, how to stop, and how to say no. Robots that build robots have to be made the way any robot is: one measured unit at a time, starting now.

22.1 Now

Desk is the current stage (Section 13). The next three results, each published whatever it shows:

- 1 The first built **lymb**'s reach, hold and sag, against its model (Table 15).
- 2 **rana**'s first dyno run, against its predicted 36 N·m peak.
- 3 The preregistration of the K1 body-schema fit, before the fit runs.

22.2 Asks

Early buyers for the **lymb** and the **duet** who will share measured data. A municipal partner for a one-city, one-task litter pilot. Collaborators who want to check the numbers or build the layers. Write to info@muni.works.

APPENDIX A

Gates

Every stage closes on gates of the same form: a number, a unit, a method and a pass condition.

TABLE 13 Gates, stage by stage²⁶. Proposed targets, to be confirmed by the author. The year range for Desk is an aspiration, not a commitment.

STAGE	GATE: NUMBER, METHOD, PASS CONDITION	EUM, ORDER
Desk (yrs 1–3)	pond revenue ≥ \$10,000 a month for 3 months running, with no outside money (proposed). ≥ 10 lymbbs sold to outside buyers at ≥ 40% gross margin, after materials, labor, test and warranty (invoices, trailing 12 months). Tool-point error ≤ 2.0 mm RMS on every lymb shipped: 50 fixed poses, unloaded and at 1.0 kg, probe against a calibrated plate, printed on its receipt. rana ≥ 1.2 × RS04 in continuous stall torque per kg or per dollar, on one dyno at 25 °C, or rana is dropped	0 W → 10 kW
Shop	A duet at a customer’s table completes a named task (for example, folding 50 garments) at ≥ 95% success. A duet assembles a lymb that passes the Desk test in ≤ 30 min of human touch time. A duet builds a duet that passes its receipts at ≤ 20% bought cost, with bought share falling over 3 generations and a commanded shutdown drill passed in each	100 kW
Road	Preregistered hardware body-schema fit: held-out joint-position error ≤ 0.8 × persistence at 0.1 s and 0.5 s, and fitted link masses within ±5% of weighed parts. Rose walks 100 m on level pavement carrying two lymbbs, 0 falls in 5 of 5. A gradient co-designed part beats its hand-designed predecessor by ≥ 10% on its named objective	1 MW
City	The litter pilot renews for ≥ 12 months after a ≥ 3-month trial, with 0 contact incidents. A city renews fleet and twin for ≥ 12 months. ≥ 90% of fleet repairs done by duets in the city, by receipts	10 MW

STAGE	GATE: NUMBER, METHOD, PASS CONDITION	EUM, ORDER
Home	A branch runs ≥ 1 year on collectors it made; doubling time measured over ≥ 3 generations; bought share ≤ 5% of cost	1 GW
Moon	A seed builds a duet that passes the Earth receipts, ≥ 50% of its mass from local material, with remote approval only	1 kW
Mars	A second duet with 0 human commands between landing and its receipt	1 kW
Star	None we can set now; reported as EUM and doubling time	3.8 × 10 ²⁶ W

APPENDIX B

pond numbers

A **frog** coding task makes about 40 model calls at about 20,000 tokens of context each: roughly 800,000 tokens in and 30,000 out, counting thinking. **frog**’s own logs show that 79% of what goes in is reusable cache, which pinned sessions turn into about 4.8 times less prefill; drafting from a session’s own history decodes 1.5–2.2 times faster (estimates from **frog** sessions⁵⁰). What one task costs depends on who runs the model (Table 14).

TABLE 14 What one **frog** task costs (estimates, October 2026). Prices are list or index prices and move monthly; the per-task figures are our arithmetic on the workload above.

ROUTE	PRICE	EOST PER TASK
DeepSeek V4 Flash API	\$0.14 per million tokens in, \$0.003 cached, \$0.28 out ⁵¹	≈ \$0.03
Qwen3-Coder-Next API	\$0.11 in, \$0.80 out ⁵²	≈ \$0.11
One rented H200 running tang , 30% busy	≈ \$3.50 an hour ⁵³ ; ≈ 20 GPU-seconds a task, batched	≈ \$0.06
Frontier models	what developers	≈ \$1–3

ROUTE	PRICECOST PER TASK
	compare against

APPENDIX C

Lymb numbers

TABLE 15 Rose fork v9 specification. Every number is a model output; none has been measured on a built arm. Arm-model jaws are the simpler tool in the arm’s own model, two flat jaw plates on the same EL05, and are lighter than gripper v2.

PARAMETER	VALUE	SOURCE / STATE
Degrees of freedom	6 + gripper	¹³
Base yaw, shoulder, elbow	RobStride RS04 × 3	Peak 120 N·m
Forearm roll, wrist pitch	RobStride RS02 × 2	Roll as a cartridge
Tool roll	RobStride RS00 cartridge	
Tool flange	Ø61, 6 × M3 on Ø27	Matches gripper v2
Foot	Ø230, 1/4" 5052, 4 × Ø6.6 on Ø190	M6 to the duet plate
Plate stock	1/8" (3.175 mm) 5052	Three 15 × 15 in sheets of 1/8"; one 1/4" sheet for the foot
Router parts	20 unique	DXF + nests generated
Jaw tips from shoulder	731 mm (786 with gripper v2)	Modeled
Mass	8.21 kg (arm-model jaws)	Modeled
Mass with gripper v2 + harness	8.44 kg	Modeled: arm plus gripper ⁴²
Continuous hold, straight out	1.49 kg (arm-model jaws)	Modeled; RS04 stall derated (estimate)
Continuous hold, gripper v2	1.27 kg (shoulder-limited)	Modeled; RS04 stall derated (estimate)
Sag / twist at flange	0.13 / 0.11 mm	Modeled
Elbow / wrist range	–120 to 115° / ±135°	
Interference, 209 poses	None	Swept in CAD
Build state	Ready to cut	None built yet

Open items (from the arm repository): RS02 thread depths are assumed at 6 mm and must be measured before screws are ordered. The clock-spring cassettes for the roll joints are drawn but not designed. Crush points found by the audit still need a policy, either soft limits or force limiting.

And the motors’ unpowered “anti-backdrive” damping is undocumented at the slow speed of a falling arm. It will be tested on the rig¹².

TABLE 16 Lymb margin sensitivity and break-even (estimate, from the figures above). A **lymb**’s economics are mostly the actuator’s economics.

LYMB VARIABLE	MATERIALS	MATERIALS MARGIN	NOTE
Base case	\$2,000	60%	~\$1,200 of it is six RobStride joints
Actuator cost +25%	\$2,300	54%	Tariffs, currency
Actuator cost +50%	\$2,600	48%	Supply shock
rana at \$100/joint (target)	\$1,400	~72%	Speculative
rana at \$1,086/joint (today’s prototype ²³)	~\$7,300	about –46%	Materials exceed list
Labor, test and warranty before margin reaches zero	~\$3,000/ lymb	0%	List minus materials

APPENDIX D

Gripper numbers

TABLE 17 Squeeze needed per finger⁴². $F = m g SF / (2\mu)$, $\mu = 0.5$, $SF = 2$; torque $T = 2 r F$, $r = 8$ mm, with % of EL05 rated torque in brackets.

OBJECT	MASS	SQUEEZE	EL05 TORQUE
Ø66 aluminum can, empty	15 g	0.29 N	0.005 N·m (0.3%)
500 ml bottle	0.55 kg	10.8 N	0.17 N·m (9.6%)
1/8" 5052 plate 300 × 200	0.51 kg	10.0 N	0.16 N·m (8.9%)

TABLE 18 Gripper v2 mass budget (target ≤ 500 g)⁴². The optional ISO 9409 adapter adds 52 g.

PART	MASS
Side plates 70 × 90 (2)	108.0 g
Top plate (70 × 56 + Ø61 lobes)	33.0 g
Palm plate	21.9 g
Fingers 20 × 80 (2)	27.4 g
Racks, 6061 (2)	10.2 g
Pinion Ø16 + shaft, steel	14.9 g

PART	MASS
Flat tips, 6061 (2)	17.3 g
RobStride EL05	242.0 g
6 × M3 × 10 screws	3.6 g
Camera, magnets, dowels, fasteners	29.0 g
Lightening windows	~22.3 g
Total at the flange	485.1 g

TABLE 19 Gripper force at rated and peak torque. $\mu = 0.5$, two faces, safety factor 2, rack friction ignored; with an assumed 0.8 rack efficiency, scale by 0.8 (estimate). Max lift is the gripper's alone. The arm holds 1.27 kg straight out (modeled), so the arm, not the gripper, sets payload⁴²

CASE	TORQUE	SQUEEZE / FINGER	MAX LIFT
Rated	1.8 N·m	112.5 N	5.7 kg
Peak (do not hold)	6 N·m	375 N	19.1 kg

APPENDIX E

Bill of materials

TABLE 20 Arm bill of materials summary^{13,26}. Arm materials ~\$2,000 (est.), of which ~\$1,200 is actuators; other lines are not yet broken out. The arm-repo BOM lists the gripper servo generically; the EL05 choice is from the gripper v2 design. The USB pair must stay within USB 2.0's ~1.5–2 m; at the full 2.1 m harness length this is open.

ARM (ONE)	CONTENT	COST
Actuators	RS04 × 3, RS02 × 2, RS00 × 1	~\$1,200 (est.)
Gripper actuator	EL05 × 1 ⁴²	–
Router parts	20 unique, 3 sheets 1/8" + 1 sheet 1/4" 5052	–
Turned parts	2 RS04 sleeves, 1 RS02 sleeve (6061)	–
Purchased	3 flanged plain bushings, 3 retaining rings, 88 screws	–
Printed	3 hub caps, base collar (PETG)	–
Harness	Flat CAN pair + 48 V pair, 2.1 m; palm-camera USB 2.0 pair in the same run ⁴²	–
Arm total		~\$2,000 (est.)

TABLE 21 rana-100 v0.1 prototype bill of materials²⁵. All prices are estimates, not binding quotes. The raw-material line is our estimate, not part of the BOM.

LINE ITEM	COST	NOTE
Housing set, 6061, CNC	\$231.30	Outsourced machining

LINE ITEM	COST	NOTE
Magnetic steel set, 1018, CNC	\$313.30	Outsourced machining
Gear blanks, steel, CNC	\$172.51	Teeth not yet cut; cutting ~\$150–300 (est.)
Manufactured subtotal	\$717.11	66% of total
COTS: bearings, magnets, encoder, wire, fasteners	\$103.36	
ODrive S1 bench drive	\$149.00	
Consumables	\$45.00	
Shipping (est.)	\$72.00	
Total (est.)	\$1,086.47	
Raw material in the machined parts	\$60–90	Estimate

References

1. Municipal Robotics. Robotic Systems for Sidewalk Maintenance: Why aren't robots shoveling the snow? (2025).
2. NASA. A Crescent Earth (art002e004441). <https://www.nasa.gov/image-detail/amf-art002e004441/> (2026).
3. Smil, V. *Energy Transitions: Global and National Perspectives*. (Praeger, 2017).
4. Energy Institute. Statistical Review of World Energy 2024. (2024).
5. Kardashev, N. S. Transmission of information by extraterrestrial civilizations. *Soviet Astronomy* **8**, 217–221 (1964).
6. Sagan, C. *The Cosmic Connection: An Extraterrestrial Perspective*. (Doubleday, 1973).
7. International Astronomical Union. Resolution B3 on recommended nominal conversion constants for selected solar and planetary properties. (2015).
8. CSIS Aerospace Security Project. Space launch to low Earth orbit: how much does it cost? (2022).
9. Dodd, M. S., Papineau, D., Grenne, T. & others. Evidence for early life in Earth's oldest hydrothermal vent precipitates. *Nature* **543**, 60–64 (2017).
10. Daeschler, E. B., Shubin, N. H. & Jenkins, F. A. A Devonian tetrapod-like fish and the evolution of the tetrapod body plan. *Nature* **440**, 757–763 (2006).

11. Harmand, S., Lewis, J. E., Feibel, C. S. & others. 3.3-million-year-old stone tools from Lomekwi 3, West Turkana, Kenya. *Nature* **521**, 310–315 (2015).
12. RobStride Dynamics. RobStride RS00, RS02, RS04, RS06 and EL05 actuator manuals and outline drawings. (2026).
13. Municipal Robotics. **rose** arm, fork concept v9: README, rose_fork.py, INTERFACES.md and fab/BOM.md. (2026).
14. Pedersen, C. **IPSE**: Interoceptive Predictive Self-Emergence. Design brief and repository status. (2026).
15. Tobin, J. *et al.* Domain Randomization for Transferring Deep Neural Networks from Simulation to the Real World. in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2017).
16. Unitree Robotics. Unitree G1 Humanoid Agent. (2024).
17. Booster Robotics. Booster K1 humanoid robot. (2025).
18. Gleditsch, N. P., Wallensteen, P., Eriksson, M., Sollenberg, M. & Strand, H. Armed conflict 1946–2001: a new dataset. *Journal of Peace Research* **39**, 615–637 (2002).
19. NASA/JPL-Caltech. Pale Blue Dot (PIA00452). <https://photojournal.jpl.nasa.gov/catalog/PIA00452> (1990).
20. Encyclopaedia Britannica. 2026 Iran war. (2026).
21. GlobalSecurity.org. Iran War 2026: Day 229 update, 14 October 2026. (2026).
22. CBS News. Massive trove of Epstein files released by DOJ, including 3 million documents and photos. (2026).
23. UN Office of the High Commissioner for Human Rights. Flawed 'Epstein Files' disclosures undermine accountability for grave crimes against women and girls. (2026).
24. Federation of American Scientists. Status of World Nuclear Forces 2025. (2025).
25. Municipal Robotics. **rana**-100 v0.1 prototype: orderable bill of materials. (2026).
26. Municipal Robotics. Municipal Robotics: Vision. (2026).
27. Ljung, L. *System Identification: Theory for the User*. (Prentice Hall, 1999).
28. Open X-Embodiment Collaboration. Open X-Embodiment: Robotic Learning Datasets and RT-X Models. <https://arxiv.org/abs/2310.08864> (2023).
29. Pedersen, C. **pond**: founding strategy document. (2026).
30. PTC. Onshape: cloud-native CAD. <https://www.onshape.com/> (2026).
31. Open Robotics. URDF: Unified Robot Description Format. <https://wiki.ros.org/urdf> (2026).
32. NVIDIA. Isaac Sim. <https://developer.nvidia.com/isaac/sim> (2026).
33. Todorov, E., Erez, T. & Tassa, Y. MuJoCo: A physics engine for model-based control. in *IEEE/RSJ International Conference on Intelligent Robots and Systems* (2012).
34. Macenski, S., Foote, T., Gerkey, B., Lalancette, C. & Woodall, W. Robot Operating System 2: Design, architecture, and uses in the wild. *Science Robotics* **7**, (2022).
35. Municipal Robotics. Vision decisions log. (2026).
36. Municipal Robotics. Battery mockup: **rose** A/B 13S2P 21700 **pack** and passive swap dock. (2026).
37. Molicel. INR-21700-P50B technical data sheet. (2023).
38. Municipal Robotics. Compute mockup: 10-inch mini rack v2 and the blade standard. (2026).
39. ISO. ISO/TS 15066:2016 Robots and robotic devices: Collaborative robots. (2016).
40. ISO. ISO 13482:2014 Robots and robotic devices: Safety requirements for personal care robots. (2014).
41. Drexler, K. E. *Engines of Creation: The Coming Era of Nanotechnology*. (Anchor Books, 1986).
42. Municipal Robotics. Gripper v2: EL05 rack-and-pinion parallel gripper, README and numbers.txt. (2026).
43. Pedersen, C. **ipse** deploy notes: K1 access and onboard computers. (2026).
44. Pedersen, C. **rose**: the humanoid line (mini, full, mega). (2026).

45. Butlin, P., Long, R. & others. Consciousness in Artificial Intelligence: Insights from the Science of Consciousness. (2023).
46. U.S. Census Bureau. QuickFacts: Lakewood city, Ohio. <https://www.census.gov/quickfacts/lakewoodcityohio> (2024).
47. City of Lakewood, Ohio. Active Living and Walkability. <https://www.lakewoodoh.gov/active-living/> (2024).
48. Dyson, F. J. Search for artificial stellar sources of infrared radiation. *Science* **131**, 1667–1668 (1960).
49. Armstrong, S. & Sandberg, A. Eternity in six hours: intergalactic spreading of intelligent life and sharpening the Fermi paradox. *Acta Astronautica* **89**, 1–13 (2013).
50. Municipal Robotics. **frog**: roadmap and session measurements. (2026).
51. Morph. Best open-source coding model 2026: Kimi K3 vs GLM-5.2 vs DeepSeek V4 vs Qwen3. <https://www.morphllm.com/best-open-source-coding-model-2026> (2026).
52. TheRouter.ai. Qwen3 vs DeepSeek for coding: API pricing, benchmarks, and routing tradeoffs. <https://therouter.ai/blog/qwen3-vs-deepseek-coding-comparison/> (2026).
53. Kotra, V. GPU cloud pricing, October 2026: what an H100, H200, B200 or B300 costs per hour. https://dev.to/viswaretas_kotra/gpu-cloud-pricing-october-2026-what-an-h100-h200-b200-or-b300-costs-per-hour-5a25 (2026).