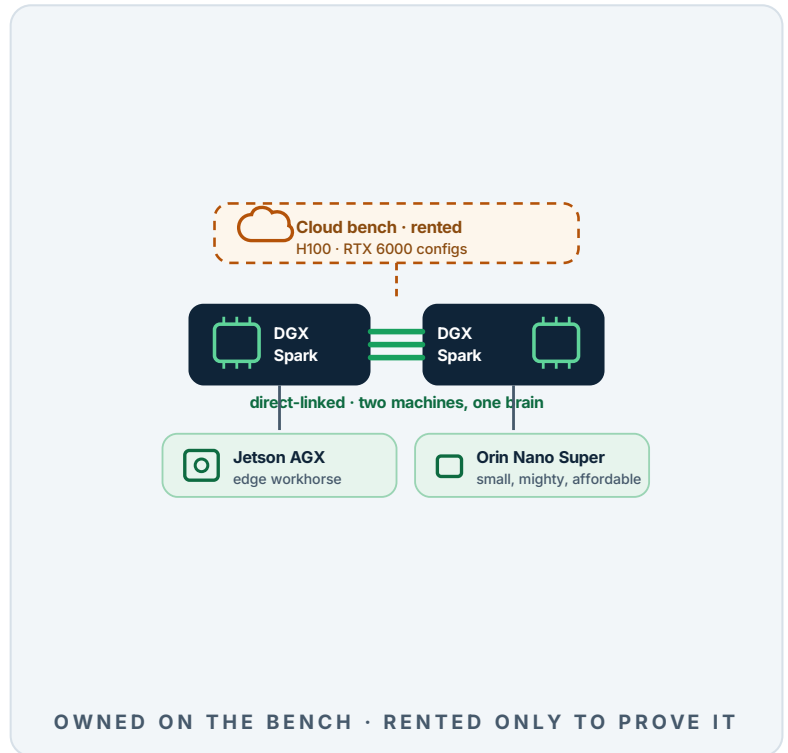


WHAT MAKES US DIFFERENT

Why trust advice that's never been plugged in?

Our recommendations come off a workbench, not a slide. We run the machines we suggest — palm-sized board to desk-side supercomputer pair — and rent cloud GPUs only to prove what's worth renting.



THE IDEA, IN PLAIN TERMS

Every machine you might buy — or rent — already lives in our lab.

When we say a \$249 board can handle your pilot, it's because we ran your kind of task on one. When we say two desk-side machines can replace a cloud bill, it's because ours are wired together right now, doing exactly that.

<p>1 The desk-side pair</p> <p>2x NVIDIA DGX SPARK · DIRECT-LINKED</p> <ul style="list-style-type: none"> Two AI supercomputers joined by a dedicated high-speed link. They pool their memory — models too big for one machine span both. Data-center-class work, on a desk, behind your own door. 	<p>2 The edge workhorse</p> <p>NVIDIA JETSON AGX</p> <ul style="list-style-type: none"> Serious AI in a box that lives next to the work — a floor, a vehicle, a bench. Runs around the clock on a fraction of the power. Answers locally even when the internet doesn't. 	<p>3 The honest little board</p> <p>NVIDIA JETSON ORIN NANO SUPER</p> <ul style="list-style-type: none"> Palm-sized, roughly the price of an office chair. Our proving ground: if your task runs here, it runs anywhere. The cheapest possible way to find out what you actually need. 	<p>4 The cloud bench</p> <p>H100 · RTX 6000 ON DEMAND</p> <ul style="list-style-type: none"> We spin up the exact configurations you'd be paying for. Your workload, timed and priced on each — before you commit. Rented by the hour to prove a point, then shut down.
---	---	---	--

THE SKILL UNDERNEATH IT ALL

Machines that talk to each other.

The hard part isn't any single computer — it's making them cross-function. We practice that daily: one codebase, deployed seamlessly from board to pair to cloud.

ONE CODEBASE · FOUR TARGETS

- Orin Nano Super · the pilot
- Jetson AGX · the edge
- DGX Spark pair · the desk
- H100 / RTX 6000 · the burst

WHY IT MATTERS TO YOU

Advice you can hold in your hand beats advice you have to take on faith.

Anyone can quote a spec sheet. We hand you measurements: how fast your task ran, what it cost per hour, and on which machine — from a \$249 board to an H100 — so the decision rests on evidence instead of hope.

THE USUAL WAY

Advice from a slide deck

- Recommendations copied from vendor brochures.
- "It should work" — discovered after the purchase order.
- Cloud bills sized by guesswork, corrected by invoice.
- Engineers meet the hardware for the first time in production.

OUR WAY

Advice from a working lab


- Your workload runs on our bench before you spend a dollar.
- Cross-machine deployment practiced daily, not improvised.
- Cloud configs — H100 vs RTX 6000 — tested head-to-head.
- You get the numbers: speed, cost, and the right-sized pick.

On the bench today: **DGX Spark x2** direct-linked · **Jetson AGX** · **Orin Nano Super** · **H100 & RTX 6000** in the cloud.

THE PROVING PATH

Start small. Step up only when the numbers say so.

Every engagement climbs the same ladder: prove the idea on the little board, harden it on the edge box, scale it on the linked pair, and burst to rented H100s or RTX 6000s only for the work that truly earns it.



Tested before you buy

A written hardware-fit report: your task, timed and priced across every tier of the lab — so the purchase order is the last step, not the first.

HELPFUL FOR LEADERS

Code that moves without drama

We package your software so the same build deploys to the board, the desk, and the cloud — checked automatically before every release.

HELPFUL FOR EVERYONE

Engineers who speak every machine

We train your team on cross-machine work — linking, load-sharing, and deploying — the skill that turns hardware into capability.

HELPFUL FOR YOUR ENGINEERS

A pilot that costs lunch, not a quarter

Proving an idea on the little board costs almost nothing. Most bad ideas die cheaply there — and the good ones climb with confidence.

HELPFUL DAY TO DAY

Want to see your workload on real hardware?

Bring us one real task — we'll run it across the lab, Nano to H100, and show you the numbers. No pressure.

EMAIL brent@llmadvisor.ai
WEB www.llmadvisor.ai
PHONE 617.595.8092