

# Spare computing power. Useful AI.

GPUs are the powerful chips that run AI. We put their spare capacity to work.

YOU / 01

What can I use LOWTIDE for?

...




Idle capacity

LOWTIDE / CAPACITY

## Run AI without owning the hardware.

Use an open AI model to summarize, sort, or pull information from text. LOWTIDE is being built to run that work on spare computing power.

SPARE POWER → USEFUL WORK



Find power



See price



Set limit



Run AI

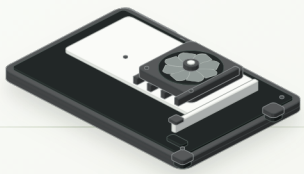


Keep receipt

THE RESEARCH CASE

# Why build on spare GPUs?

Paying for computing power and putting it to useful work are different things. LOWTIDE starts with a practical question: can spare, authorized capacity run a suitable open model at a lower total cost?



GPU capacity

Hardware illustration

MICROSOFT TRAINING CLUSTER · 2019

## 52.3%

Mean utilization of allocated GPUs

In a two-month Microsoft study of deep-learning training, allocated GPUs averaged 52.32% utilization. This shows an efficiency gap in one system. Finding capacity that can accept useful work is a separate problem. [1]

One training-cluster study. This is not a measure of all GPUs or of capacity available for resale.

01

### Unused headroom exists.

The Philly study found substantial underuse even on GPUs already allocated to training jobs. The next question is whether that headroom can be shared safely and usefully. [1]

PHILLY
USENIX ATC 2019

02

### Cheap capacity can disappear.

SpotServe studies language models on interruptible cloud machines. Its experiments show a trade-off: lower compute costs can come with more time waiting for a response. [2]

SPOTSERVE
ASPLOS 2024

03

### Getting a model ready takes work.

ServerlessLLM tackles the time spent loading model weights. Faster startup can make flexible capacity more practical, but it is only one part of delivering a useful answer. [3]

SERVERLESSLLM
OSDI 2024

AN ILLUSTRATIVE WORKFLOW

## A product catalog that can wait.

An example, not a customer result.

A team needs to classify a backlog of product descriptions. It checks an open model against reviewed examples, then asks for a current quote and authorizes a spending cap. If no slot is free, its application can try again later.

| 01 | 02 | 03 |
| --- | --- | --- |
| Check the quality | Cap the request | Keep the receipt |

The task can wait; the answer still has to meet the team's quality bar. This is a candidate for the proposed pilot, with no measured LOWTIDE saving implied.

LOWTIDE · PROPOSED PILOT

## How that shapes LOWTIDE.

Useful output. Reliable requests. A sustainable cost. The pilot has to prove all three together.

Use capacity that is genuinely available. The plan requires permission to serve third-party work and an available request slot. A low GPU activity reading alone is not enough. Product plan ↗

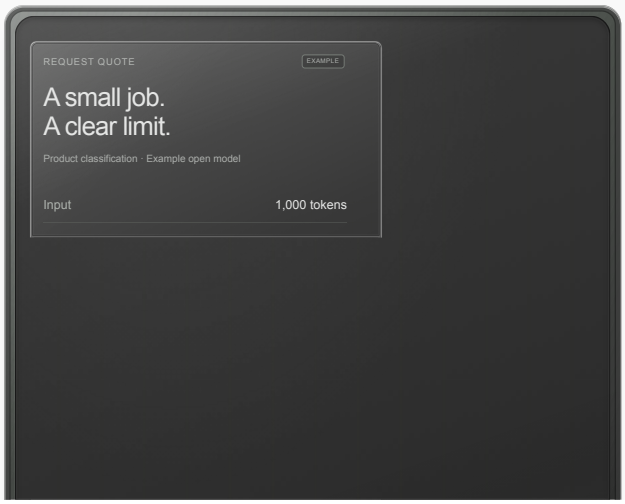
Price the service, including the work that fails. Proposed prices respond to occupied serving slots, with a minimum that covers the measured cost of compute, operations, and failed work. Lower activity does not guarantee a discount. Product plan ↗

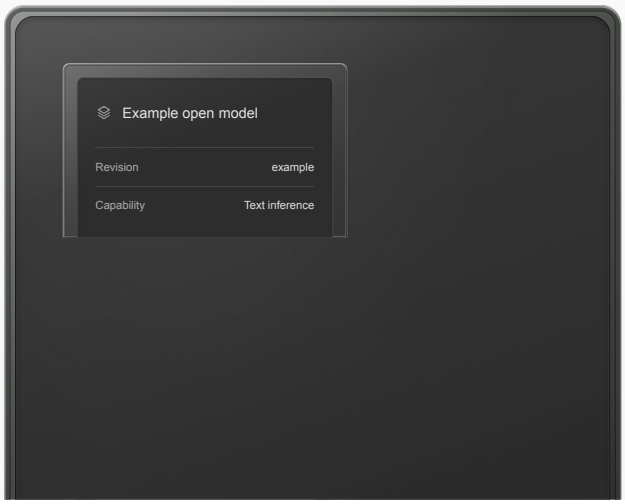
Measure before making a promise. Quality, successful requests, availability, and total cost must pass the pilot's checks. Showing a quote and a spending cap does not establish a price advantage.

Read the planned pilot and its limits ↗

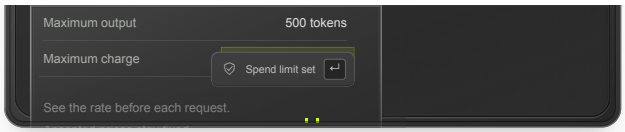
Read research notes 7 primary sources +

# Useful compute, in every form.

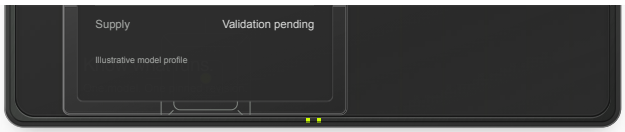


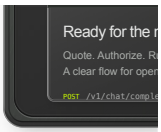



**Quotes.** Know your maximum before you run.




**Models.** Choose a revision that fits your work.

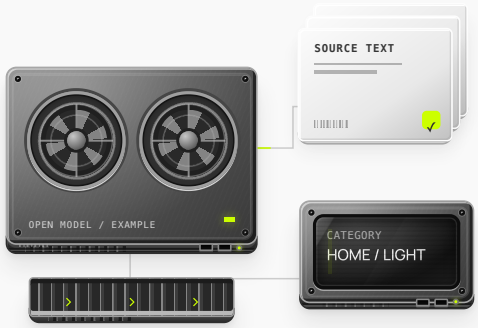



**Discovery.** Keep

18 WORKFLOWS / ONE CONNECTED WORLD

Pause motion

STRUCTURE / 03 FLOWS




SOURCE → FIELDS 01

WRITING



01 / STRUCTURE ILLUSTRATIVE

## Give text a structure.

Turn source material into fields you can inspect. Keep unknowns explicit.

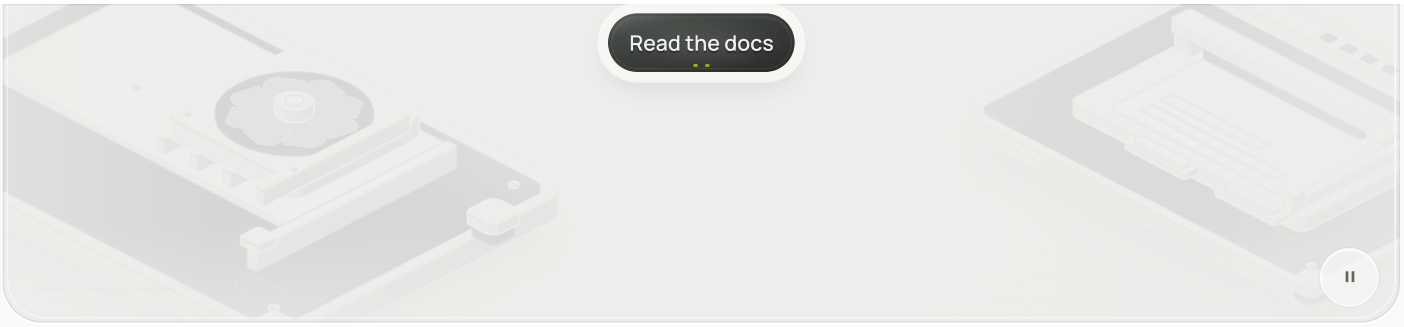
- Product classification
- Field extraction
- Record normalization

No request sent. Review every output.

01 Structure | 02 Writing | 03 Operations | 04 Discovery | 05 Review | 06 Delivery

▶ Browse all 18 workflows

LOWTIDE / CAPACITY, UP CLOSE




### Product

Features

Workflows

Productivity

FAQ

Get started

### Solutions

Builders

Developers

Teams

### Resources

Why spare GPUs?

Developer flow

Articles

Console features

### Information

Docs

Whitepaper

Roadmap

For holders

Terms status

Privacy

Pricing

### Explore

Getting started

Workflows

Overview PDF

Build PDF

Billing records

LOWTIDE
Product information

©2026

LOWTIDE