

Series A - Munich, Germany

Wardline

Predicts failing GPUs and network links in training clusters and drains them before a job crashes

The problem


What's broken today

On a 4,000-GPU training run, one failing GPU, memory error or flapping network link can crash the whole job, and teams lose hours rolling back to the last checkpoint. Operators at that scale report interruptions every few hours, and the warning signs, rising ECC error counts, thermal throttling, link retries, sit unread in separate logs.

Our solution

Wardline Control

Wardline runs a lightweight agent on every node that streams hardware telemetry to a model trained on more than 9,000 recorded failures. When a node's failure risk crosses a threshold, it triggers a checkpoint, cordons the node in the scheduler and swaps in a spare, so the job carries on instead of restarting from an old checkpoint.



Incident INC-3318

 Wardline Control

Fleet

Risk

Incidents

Runbooks

Settings

Incident INC-3318

gpu-n0217, job llm-pretrain-7b-r4

Search risk...

Open RMA with vendor

Warning lead time

41 min

Checkpoint taken

12:04

Spare swapped in

12:06

Compute saved

~1,900 GPU-h

Timeline

Time	Event	Source	Detail
11:23	Risk crosses 0.80	Wardline model	ECC count +240%
11:58	Risk 0.91	Wardline model	Uncorrectable error
12:04	Checkpoint written	Training job	Step 182,400
12:06	Node cordoned	Scheduler	Spare gpu-n0509 in
12:31	GPU failed	Node agent	Confirmed by vendor

Showing 5 of 372

View all →

Insight

Without the early checkpoint, the job would have rolled back 3 h 10 m to step 176,800.

 Wardline

JR Jakob Reinhardt

Why now

The timing

Training clusters grew from hundreds to tens of thousands of GPUs in three years, and at that size hardware failure is a daily event, not an edge case.

Market

AI labs, GPU cloud providers and enterprises running clusters above 500 GPUs; roughly 1,200 such operators worldwide, managing an estimated 6M accelerators by 2027.



Wardline

How we make money

Subscription priced per monitored GPU, about US\$12 per GPU per month, with discounts above 10,000 GPUs and an on-premise edition for air-gapped sites.

Traction

Where we are

US\$5.8M

Annual recurring revenue

41,000

GPUs monitored

23

Customers

87%

Failures caught early

- US\$5.8M ARR from 23 customers, up 3.1x year on year
- 41,000 GPUs monitored across nine countries
- 87% of hardware failures caught before a job crashed



Wardline

Team

Who's building it

JR

Jakob Reinhardt
CEO and Co-founder

MH

Mira Hoffmann
CTO and Co-founder

EY

Emre Yildiz
Head of Machine Learning

MD

Margaux Dubois
VP Customer Success

The ask

What we're raising

Raising US\$28M Series B to cover 250,000 GPUs by 2028

- Add switch and storage telemetry to the risk model
- Build a US team for GPU cloud providers
- Ship an air-gapped edition for public research centres