



AI Training Platform

Maxis



Contents

1.	Overview	02
2.	Business Impact	05
3.	Core Capabilities Maxis (AI Training)	07
4.	Observability, Security & Operations	11
5.	Works with	14

1. Overview



Train models on your enterprise data across scattered xPU resources.

Development environments, training, and evaluation run in one place. Kueue-based queues and priorities allocate GPUs, so several teams share capacity in turn.

01 Scheduling and resource allocation

When GPUs run short, requests queue and start as capacity frees up. Urgent training can go first, and per-project quotas stop one team taking everything.

02 Heterogeneous xPU execution

GPUs and NPUs spread across several clusters are handled from one place. Distributed training runs across multiple GPUs.

03 An unbroken path

From an approved model → refined on your data → compared → into service. Four steps in one place, so you never switch tools between steps or lose results.

04 Compare results

Put several training runs side by side and see which one is better.

2.

Business Impact



Turn model development investment into real service results.

Maxis brings Kueue-based scheduling, workload execution, and the full development, training, and evaluation cycle onto a single platform.

01 Faster time to market

Verified job templates and one-click fine-tuning

02 Unit cost and efficiency

Kueue-based queueing and priority scheduling with per-project GPU allocation

03 Experiment reproducibility

Experiment reproduction backed by pipelines and version control

3.

Core Capabilities



Preconfigured development environments, ready with no setup

Name it, pick the resources you need, and the environment is up in minutes. Environments you are not using shut down on their own, so resources do not leak.

A development environment that starts right away

No installation or setup; you develop straight from the browser. Name it, pick the resources you need, and the environment is up in minutes.

Only what you use

Environments you are not using shut down on their own. Nothing leaks because someone left one running.



Fine-tuning that refines approved models on your enterprise data

Start from the models and datasets your company has approved. For tabular data, decide what to predict and a suitable method is found for you.

Text Generation: enterprise-grade fine-tuning

An advanced platform for fine-tuning a company's own AI models. Supports fine-tuning (Full · LoRA · QLoRA) and GKD.

Automatic training on tabular data

Build a model from spreadsheet-style tabular data. Decide what to predict and a suitable method is found for you.

Choosing a model and data

Pick an approved model from the list, add your dataset, and start training. No downloading or fixing paths every time.

Keeping the result

A finished training run is registered as a model the team shares. A record of who built which model stays in one place.

Project isolation

You see only your own project. Resources other teams create do not appear in the list, so departments and customers cannot see each other's data.



An automated pipeline from data preparation through evaluation

Steps that do not fit refuse to connect, so a badly built flow never runs.

Drawing the training flow

Connect blocks to lay out the order from data preparation through training and evaluation. Steps that do not fit refuse to connect, so a badly built flow never runs.

Run control

Only the actions available right now appear as buttons — starting, stopping and retrying never get confusing.

Runs on a schedule you set

Set a cadence such as daily or weekly and it runs on its own. No re-triggering by hand each time new data lands.



Evaluation and registration that compare models and track versions

Put several training runs side by side and see which one is better. It is always clear which version is in service.

Quality measurement

Measure how well a model performs. Latency, cost, accuracy, and task-specific metrics are scored automatically and candidates compared side by side.

Version control

Results pile up as you train again and again; each one is named and tracked. It is always clear which version is in service.

Data-driven decisions

Optimize model and prompt choices and reduce production rollout risk.

HITL evaluation workflow

Human-expert evaluation for subjective tasks

Content filters and guardrails

Automated safeguards for safety checks, policy limits, and compliance

4.

Observability, Security & Operations



Visibility, security, and operational control for enterprise AI.

See how much is being used and what training is running on one screen. Administrators adjust resource allocation from a separate console.

Usage at a glance

See how much is being used, what training is running, and what happened recently, all on one screen.

Administrator console

Administrators see who is using how much from a separate screen and adjust resource allocation.

Resource allocation

When GPUs run short, requests queue and start as capacity frees up. Urgent training can go first, and per-project quotas stop one team taking everything.

Security and authentication

Signum provides the shared foundation for access, permissions, and auditing.

5.

Works with



A model you build once keeps working across the team.

A model fine-tuned in Maxis is saved to the asset hub automatically. Another person starts from it, refines it again, compares results, and puts it into service.

01 Model fine-tuning

Pick an approved model and fine-tune it on your enterprise data.

02 Automatic asset hub registration

A finished model is saved to the asset hub automatically and is visible to the whole team. There is no separate upload step.

03 Further fine-tuning

Another person takes that model as a starting point and fine-tunes it again. Nobody starts from scratch.

04 Performance comparison

Line up the candidates you have and measure which one is better.

05 Inference deployment

Push the model you choose to a Metis inference endpoint and run it as a service.

Used together with

Signum asset hub and the shared foundation for access, permissions, and auditing · Capsis Kubernetes execution · Metis inference

CTA

Contact sales · Request a demo