

CASE STUDY

How InstaHeadshots scales AI-generated portraits with Runpod

Serverless GPU infrastructure helped the team cut cost and improve performance.

100%

performance improvement

50%

infrastructure cost reduction

1 day

migration time

The challenge

InstaHeadshots had fluctuating demand and diverse GPU needs across AI image generation workloads, which made traditional cloud infrastructure inefficient.

The Runpod approach

Runpod Serverless let the team autoscale resources, select fit-for-purpose GPUs, and pay for compute used rather than idle capacity.

The result

The migration helped InstaHeadshots improve processing performance, reduce GPU costs, and keep engineering focused on growth and product development.