



E42 gets Hyperscaler-Grade Kubernetes Performance at a Fraction of the Cost

From a stalled Kubernetes rollout to a stable, scalable AI infrastructure with AceCloud.

INDUSTRY
AI Startup

LOCATION
Pune, Maharashtra, India

OVERVIEW
vCPUs with Managed Kubernetes & Block Storage for AI Compute

Key Outcomes

95%/80%
Dev and UAT migrated

Matched
Hyperscalers
Better on throughput & RAM

3rd-Gen
Hardware available on-demand

Overview

E42 needed a Kubernetes-native, high-performance storage environment capable of supporting its growing multi-cloud AI infrastructure. Following a failed deployment with another provider that impacted performance and reliability, the company turned to AceCloud for a solution. AceCloud successfully designed and deployed a scalable, resilient storage architecture that optimized performance, simplified Kubernetes operations, and ensured the reliability needed to power E42’s AI-driven workloads across multiple cloud environments.

Cost, Kubernetes performance, and a support team that actually shows up. This combination made us move our infrastructure to AceCloud.

-Kishor Garud, Infra Head, E42.ai

Use Case

- Storage-bound AI Compute
Is heavily dependent on block storage and Kubernetes resources to support high performance compute for AI workloads
- Kubernetes-native by Design
Runs every workload on managed Kubernetes, with its entire architecture built around K8s
- 3 Clouds, 1 Strategy
Operates multi-cloud across AWS, GCP, and now AceCloud

Challenges

- Kubernetes Deployment that Didn’t Deliver
An earlier attempt to run Kubernetes workloads on a different provider surfaced persistent block storage instability and performance issues, stalling their rollout before it could scale.
- Balancing Cost Efficiency with Enterprise Grade Reliability
As an AI-native, multi-cloud organization, E42 needed a provider that could match hyperscaler performance without the cost premium attached to it.

Solutions

- AceCloud delivered managed Kubernetes infrastructure built for cost efficiency, performance, and reliability at AI scale.
- Managed Kubernetes-as-a-Service
Performative, production-ready K8s that cleared E42’s bar after POC evaluation.
 - Third-Generation Hardware, On-Demand
Newer-generation hardware delivers better RAM/CPU performance without the separate enterprise agreements required on GCP.
 - High-Performance Block Storage
Stable, high-throughput block storage that resolved the exact reliability gaps that blocked their earlier migration attempt.
 - Dedicated Human Support
Hands-on support that resolved early deployment issues quickly.

Impact Highlights

E42 gained hyperscaler-grade Kubernetes performance at a fraction of the price, backed by 21 Kubernetes nodes and active backup and disaster recovery from day one.