

FROM AN IDEA TO LIVE AI APPLICATION IN 48 HOURS



Neeraj Singh

Senior Manager – Vayu AI Cloud Product Offering

Neeraj leads Product Offering initiatives for Tata Communications' Vayu AI Cloud and Platform services.

With more than 10 years of experience in AI and cloud-native technologies, he transforms complex business challenges into scalable product and solution offerings that deliver tangible value.

He leads product go-to-market strategy, offering development, and positioning, translating technical capabilities into compelling business outcomes for enterprises, SMBs, and large organizations.

Neeraj combines customer insight, technical expertise, and commercial thinking to align product innovation with evolving market needs.

Executive summary

Consider a familiar scenario. On Monday, a business team identifies a high-value AI use case and receives leadership approval to move forward. Everyone expects development to begin immediately.

Instead, the rest of the week is spent requesting GPU access, provisioning infrastructure, installing frameworks, configuring vector databases, integrating tools, and obtaining security approvals. Before a single prompt is tested or a model is evaluated, the project has become an infrastructure exercise.

The challenge is rarely the AI idea itself. In many cases, it is not even the AI model. The real bottleneck is assembling the platform required to build, deploy, and operate AI applications. By the time the environment is ready, weeks or even months may have passed before anyone can determine whether the use case delivers business value.

This is why many enterprise AI initiatives stall before they truly begin. The common obstacles are operational rather than technical:

**Slow time to value:**

Teams must arrange infrastructure, drivers, pipelines, storage, monitoring and approvals before development begins.

**Rigid infrastructure and scarce GPUs:**

Capacity is expensive, slow to procure and often underused between projects.

**Fragmented tools:**

Notebooks, pipelines, registries, deployment and monitoring sit across disconnected products, creating friction at every hand-off.

**Governance arriving late:**

Security, compliance, and responsible AI controls are added after experimentation, delaying production.

**Unpredictable cost:**

Idle GPUs and poorly optimised workloads make cloud spending difficult to forecast.

These are not AI model problems. They are assembly problems, and they can consume months before a team measures a business outcome.

One platform for the entire AI lifecycle

Tata Communications' Vayu AI Studio, built on Vayu AI Cloud, removes much of that assembly work. The fully managed, enterprise-ready environment supports the lifecycle in one place:

**Discover:**

Ready-to-consume AI Model hub to evaluate and choose preferred AI model.

**Build:**

Preconfigured notebooks and IDEs, Vector Databases, distributed training, experiment tracking and versioning.

**Deploy and scale:**

Managed production endpoints, staged rollouts and real-time performance monitoring.

**Govern:**

Role-based access, audit logs, responsible-AI guardrails and policy controls, with data and metadata hosted and managed in India.

The real test was whether first-time users could turn ideas into working applications under pressure.

A 48-hour test with first-time users

The first Tata Communications' Vayu AI Studio Hackathon drew 304 students across 98 teams from three Indian university campuses. Fifteen shortlisted teams received documentation and a support channel, then had 48 hours to build and deploy a working AI application on Vayu AI Cloud.

They were students, not specialist AI engineers. None had used the platform before or had an MLOps specialist to rely on. All 15 teams delivered a live application within 48 hours.

What the teams built



IT operations

An assistant that reads alerts, gathers logs, suggests likely causes and asks for human approval before changing a live system. It used Tata Communications' Vayu AI Studio workspaces and the platform's model-serving and monitoring services.



Official documents

Retrieval applications that answered questions on GST circulars, laws, welfare schemes and farm leases, cited the supporting page or clause, and declined unsupported answers. Several worked in Hindi and English.



Greenhouse farming

Sensor-driven irrigation that recommends when crops need water, with a safety layer that can stop the pump and explain why.



Manufacturing quality

A computer-vision application that marks suspected circuit-board defects, applies factory rules and translates findings into batch cost.



Physical infrastructure

Early-warning systems for EV battery depots, railway tracks and data-centre cooling, combining sensor data, predictive models and manual override.

What the results show

In two days, first-time users ran language models, computer vision, live sensor streams, retrieval systems and classical machine learning on one platform. They did not need a separate integration project because versioning, registry, model serving and monitoring were already available.

The applications also used practical safeguards: Human approval before consequential changes, safety layers that could overrule a model, and systems that refused to answer without evidence. Multilingual applications ran on infrastructure hosted in India.



What the numbers say

100% uptime

for 14 of 15 teams. One infrastructure issue was handled by moving the team to a standby cluster without loss of work.

About 20 support queries per team,

with roughly 90% resolved in 5 to 10 minutes.

4.2 out of 5

overall satisfaction, with 12 of 15 teams rating the platform 4 or 5.

4.7 out of 5

for product support.

94%

would recommend Tata Communications' Vayu AI Studio to a peer.

Teams particularly valued the documentation and developer-friendly APIs, which helped them spend their time building rather than deciphering the platform.

So what is holding your team back?

Fifteen teams with no prior platform exposure and no MLOps specialist went from a blank page to a deployed application in 48 hours. Most organisations start with additional advantages, including domain knowledge and data.

The constraint is often not the ability of the team. It is the platform they are given and the assembly work between an idea and the first working prototype. A startup can test before funding infrastructure. An SMB can build without hiring a platform team. An enterprise can prove value before committing to a long platform programme, while keeping governance, sovereignty and auditability in place from day one.

If 48 hours is enough to learn whether an idea works, the cost of trying has changed. So has the cost of waiting.

Start building

Explore Tata Communications' Vayu AI Studio with complimentary credits and see what your team can build, deploy and scale in a few days, with enterprise governance and sovereignty built in.

Explore Tata Communications' Vayu AI Studio

vayu | AI
Cloud

For more information, visit us at www.tatacommunications.com

CONTACT



©2026 Tata Communications. All Rights Reserved. TATA COMMUNICATIONS and TATA are trademarks of Tata Sons Limited in certain countries.