

Tata Communications Vayu AI Cloud

AI STUDIO

Build Smarter. Deploy Faster. Govern Better

Artificial Intelligence is no longer a futuristic vision-it's a **strategic business imperative.** As the **second most discussed technology topic** among global leaders after cybersecurity, AI holds the potential to reshape industries. However, despite 86% of enterprises initiating AI adoption, only a handful have successfully scaled AI across their organisations.

The challenge lies not in the ambition but in the **execution**. Seamless **data integration**, **governance**, and accessibility form the backbone of Al success-yet 80% of Al project time is spent on **data preparation**, and more than half of enterprises grapple with **poor data quality** and **real-time processing bottlenecks**. Additionally, **talent shortages** in areas like ML pipeline management and distributed compute optimisation further stall progress.

Scaling AI profitably requires a **paradigm shift**-moving from resource-intensive pilots to **automation-driven platforms** that bridge both **data and talent gaps**. With **AI-native tools, managed services**, and a clear framework to link AI initiatives with **business outcomes**, enterprises can unlock the full potential of AI-turning pilots into enterprise-wide impact.





Al Studio: core capabilities

Tata Communications Vayu AI Studio offers an end-to-end platform with tools for data management, experimentation, training, fine-tunning and deployment-designed for speed, security, and efficiency. It comprises an Al Workbench that provides a fully managed, customisable AI workspace with pre-configured environments and AI copilot for seamless coding, troubleshooting, and deployment. Al supermarket offers access to open-source and premium models like Llama, DeepSeek, and Mistral, with Model-as-a-Service (MaaS) enabling easy API-based AI adoption. Serverless AI eliminates infrastructure complexities with on-demand scalability, supporting Training-as-a-Service, RLHF, and inference.

Al Studio



Al workbench



Al supermarket



Serverless Al



Data management



(XX) MLOps/GenAlOps



Responsible Al

Let's explore each layer in depth



Al Workbench

Al Workbench revolutionises Al development by eliminating manual framework setup and dependency management. With pre-built workspaces and an intuitive conversational Copilot, one can achieve:



Faster time-to-market:

Instant project launches, and automated workflows reduce development time.



Increased productivity:

Al-powered automation streamlines coding, troubleshooting, and deployment.



Simplified AI and ML development:

No-code interfaces and visual tools remove technical barriers for non-developers.



Optimised data processing:

Pre-built ETL pipelines accelerate data transformation and integration.



Reduced setup overhead:

Pre-configured workspaces eliminate manual environment setup and configuration with ready-to-use development environments.



Bridging skill gaps:

Al-driven tools empower SMEs and data scientists to build AI applications without deep coding expertise.







💹 Al Supermarket

Many platforms fall short by lacking **integrated, ongoing managed services,** slowing down development and deployment. Al Studio overcomes these gaps with an end-to-end "Al Supermarket" -offering robust managed services and continuous support for seamless Al innovation.

Model Garden:



Flexibility and choice:

Leverage open-source AI models like Llama, Mistral, and DeepSeek with a built-in playground to validate models using your own data.



Enterprise-grade performance:

Gain premium access to third-party models such as Sarvam and Corover, backed by SLAs for reliability and performance.



Seamless AI adoption:

Accelerate AI integration with Model-aaS—offering pre-trained models via API integrations across multiple providers like OpenAI, Llama, and more for diverse use cases.



Faster innovation:

Simplify complex AI workflows like RAG (Retrieval-Augmented Generation) or direct inference without heavy infrastructure investments.

Pre-trained APIs:

Custom-built APIs designed for smooth integration into user applications to simplify tasks like text generation, summarisation and classification.



NIMs:

A set of easy-to-use microservices designed to accelerate the deployment of foundation models. Enables developers to focus on application logic while leaving complex integrations to microservices.

Embedding support:

Ready-to-use embeddings for text and image-based tasks. Simplifies the development of applications requiring advanced understanding of text or image data.

Managed AI tools:

A one-click solution to all leading AI frameworks and Vector Databases with comprehensive managed services for seamless integration and performance.





🗷 Serverless Al

Scaling AI workloads effectively while managing costs, infrastructure, and model complexity poses a major challenge for organisations. Traditional methods demand extensive setup, resource provisioning, and ongoing maintenance-slowing agility and time-to-market.

Serverless AI transforms AI infrastructure with on-demand scalability, eliminating manual resource management. Beyond typical serverless benefits, it offers advanced AI features for a seamless, holistic experience.



Effortless model training:

Training-aaS provides fully managed, auto-scaling environments, allowing seamless model training without infrastructure complexities.



Human-guided precision:

Reinforcement Learning with Human Feedback (RLHF) improves model accuracy by aligning outputs with human preferences through interactive training loops.



Active learning:

Enhance model performance by prioritising the most valuable data for annotation and iteratively refining model accuracy.



Seamless AI deployment:

Inferencing-aaS deploys high-performing inference pipelines with low latency, auto-scaling, and seamless integration with downstream applications.



Optimised model performance:

Built-in hyperparameter tuning uses intelligent algorithms to refine model performance and minimise trial and error.



Pretraining support:

Kickstart custom AI projects with pre-trained models and datasets, enabling rapid prototyping and deployment.

With these comprehensive capabilities, Serverless AI accelerates the AI lifecycle while keeping costs under control and providing flexibility for diverse workloads.







Data Management

Fragmented datasets, inconsistent formats, and inefficient feature engineering pipelines often hinder scalable AI solutions. Our Data Management suite enables organisations to unlock the full potential of their data through seamlessly integrated, AI-driven data management systems. Addressing every stage of the data lifecycle, our platform offers:



Simplified dataset management:

Create, version, and manage datasets across environments to ensure consistency and accessibility.



Effortless data preparation:

With Data Wrangling, clean, transform, and prepare raw data with no-code data wrangling and automated tools.



Data chunking:

Leverage advanced chunking techniques to break large datasets into manageable pieces, enabling faster training and optimised storage utilisation.



Centralised feature sharing:

Store and share curated features through a feature store, ensuring consistency across training and inference workflows.

By integrating these tools, organisations can streamline data preparation, improve model accuracy, and maximise the impact of their Al initiatives.

MLOps/GenAlOps

As Retrieval-Augmented Generation (RAG) gains traction for deploying GenAl solutions, organisations often struggle with complex setup and fine-tuning. RAG-aaS simplifies this journey by enabling customers to select models, experiment with prompts, and integrate social guardrails—all through a single, intuitive interface at runtime. With built-in explainability frameworks like SHAP, it ensures responsible Al practices, offering transparency and accountability. Businesses can deploy GenAl solutions faster, with greater flexibility and ethical oversight.



Customisable model selection:

Choose the right model for your GenAI use case with integrated version management and model registry



Optimised performance:

Refine prompts through prompt engineering for more accurate and efficient outcomes.



Transparency and accountability:

Real-time explainability with built-in SHAP framework ensures clear, ethical decision-making.



Seamless deployment:

One-click integration streamlines GenAl deployment, enhancing speed and operational efficiency.







Responsible Al

As AI adoption accelerates, organisations face challenges around ethical use, security vulnerabilities, regulatory compliance, and accountability. Without robust safeguards, Al solutions risk bias, misuse, and non-compliance-eroding trust and business value.

Our Responsible AI framework embeds accountability, security, and transparency across every stage of the Al lifecycle.



Model security and privacy:

Safeguard models from adversarial attacks, ensure data privacy during training and inference, and encrypt sensitive workflows.



Ethical AI frameworks:

Mitigate bias and harmful outcomes with pre-defined social guardrails.



Traceability:

Maintain comprehensive audit logs for datasets, models, and decisions to ensure end-to-end traceability and compliance with regulations.



Proactive risk management:

Proactively identify and mitigate risks related to AI bias, misuse, and unintended consequences through real-time monitoring and alerts.



Runtime security:

Secure production environments with robust runtime monitoring and defence mechanisms to prevent unauthorised access and anomalies.

By prioritising security, transparency, and ethical Al usage, Responsible Al ensures that organisations can deploy AI systems confidently and responsibly.









