Solving for Enterprises with Generative AI

Generative AI is a cutting-edge technology with the potential to revolutionize business operations profoundly. It is being adopted increasingly across multiple industries to create personalized experiences, improve customer service, enhance productivity, accelerate innovation, and generate new business models.

Canvas.GEAR - An Entereprise-grade Generative AI platform

Canvas.GEAR, driven by Large Language Models and integral AI frameworks, has been designed to address distinct business challenges through its intuitive interface. The platform enables democratization of AI capabilities helping business users curate and train enterprise data securely to power diversified industry solutions. It provides secure, sustainable and engaging user experiences through context isolation, content moderation, and responsible AI practices.

Benefits ▶

More than 30 industry-specific trained models enabling faster go-to-market

20% to 40% cost savings through better governance and FinOps

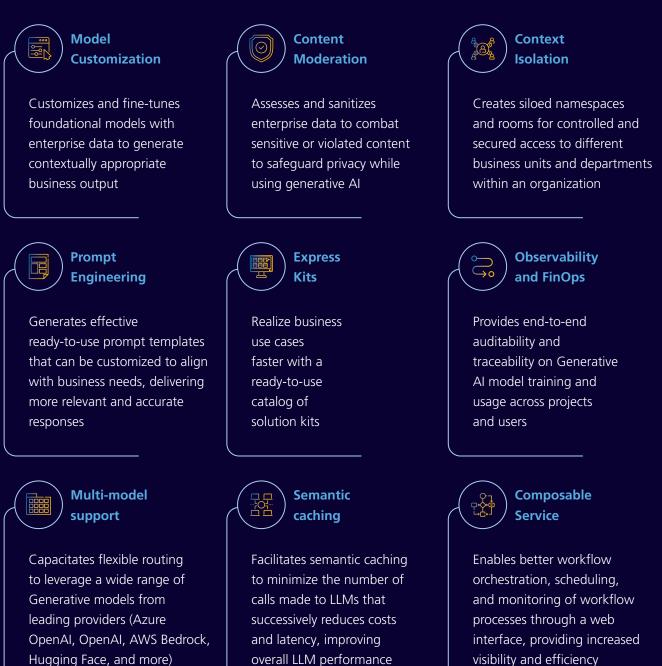Prompt engineering, enabling better user experiences

Express kits accelerating use-case enablement

Why LTIMindtree?

LTIMindtree's capabilities in global technology consulting and digital engineering are known to reimagine enterprise businesses across industries. With the tremendous surge in Generative AI, LTIMindtree is building enterprise frameworks to drive increasingly innovative, sustainable, and competitive offerings. Additionally, our strategic foresight and dedicated AI/ML experts team have helped us unveil an array of industry solutions withgenerative AI. We are among the first few to get early access to hyperscaler generative AI foundation models to build platform-based, ready-to-use, industry-specific generative AI solutions.

Salient Features

Model Customization

Customizes and fine-tunes foundational models with enterprise data to generate contextually appropriate business output

Content Moderation

Assesses and sanitizes enterprise data to combat sensitive or violated content to safeguard privacy while using generative AI

Context Isolation

Creates siloed namespaces and rooms for controlled and secured access to different business units and departments within an organization

Prompt Engineering

Generates effective ready-to-use prompt templates that can be customized to align with business needs, delivering more relevant and accurate responses

Express Kits

Realize business use cases faster with a ready-to-use catalog of solution kits

Observability and FinOps

Provides end-to-end auditability and traceability on Generative AI model training and usage across projects and users

Multi-model support

Capacitates flexible routing to leverage a wide range of Generative models from leading providers (Azure OpenAI, OpenAI, AWS Bedrock, Hugging Face, and more)

Semantic caching

Facilitates semantic caching to minimize the number of calls made to LLMs that successively reduces costs and latency, improving overall LLM performance

Composable Service

Enables better workflow orchestration, scheduling, and monitoring of workflow processes through a web interface, providing increased visibility and efficiency

Canvas.Gear Personas

Role-based access control for creation, training, consumption, and administration of content with –

CREATOR Persona

- Provides end-to-end Gen AI solutions (Setup to use case building & execution)

- Enables enterprise data training as a self-service for business users

ADMINISTRATOR Persona

- Creates isolated namespaces and manages role-based access controls for different personas

- Facilitates better governance through monitoring, rate limiting, and fair usage control

END-USER Persona

- Quick onboarding through pre-built chat applications

- Enables API integrations for seamless chat, search, code, video, and documentation solutions

---

Canvas.Gear is an enterprise-grade Generative AI platform designed to curate and train industry-specific data to build, manage, and deploy domain-specific business solutions.

- -