NOVARTIS

AI Engineer

Job ID
REQ-10093342

十一月 24, 2026

India

摘要

#LI-Hybrid

Location: Hyderabad, India

About The Role:
This role is an opportunity to turn the Corporate Affairs AI strategy into working solutions that deliver measurable value. The AI Engineer will design, build, test, deploy, and optimize enterprise Generative AI and Agentic AI applications that improve how teams generate insights, create content, and support strategic communications and engagement. Working closely with the Director of AI Enablement, product management, data engineering, architecture, and business stakeholders, the role will convert prioritized use cases into reliable, secure, and scalable solutions that are practical for day-to-day use.

The AI Engineer will be accountable for the end-to-end technical delivery and production performance of assigned AI solutions. Success will be measured by the ability to move use cases from experimentation into dependable production applications with strong user experience, sustained performance, appropriate evaluation controls, and clear business impact.

About the Role

Key Responsibilities:

- **Design, build, configure, test, deploy, and support Generative AI and Agentic AI solutions** for priority Corporate Affairs use cases, with end-to-end ownership from experimentation and technical validation through production deployment and ongoing support.
- **Build production-oriented Retrieval-Augmented Generation (RAG) solutions**, covering ingestion, chunking, embeddings, metadata filtering, vector or hybrid retrieval, reranking, prompt construction, grounding, and response generation.
- **Design and implement AI agents and stateful workflows using LangChain and LangGraph**, incorporating tool calling, orchestration, memory/state, human-in-the-loop, exception handling, and observability patterns.
- **Develop and integrate Model Context Protocol (MCP) clients and servers** encompassing tools, resources, transport mechanisms, authentication, security, and enterprise-system connectivity, while implementing enterprise integration patterns using REST APIs, API Gateways, authentication and authorization, synchronous/asynchronous communication, and resilient error handling.
- **Establish robust AI evaluation and observability approaches** using LangSmith, Ragas, or equivalent frameworks to measure retrieval quality, answer relevance, faithfulness, groundedness, hallucination, latency, and regression.
- **Develop backend services, reusable components, APIs, and user interfaces**, using Python and, where appropriate, JavaScript/TypeScript or Java for backend development, and React.js, Next.js, or comparable modern web frameworks for front-end development.
- **Develop and deploy solutions on AWS**, integrating relevant cloud, security, logging, data, and application services and applying Amazon Bedrock where suitable; work with structured and unstructured enterprise data across platforms such as Amazon S3, Databricks, and Snowflake.
- **Use Python and Jupyter Notebook for rapid experimentation, prototyping, evaluation, and technical validation**, maintaining a clear path to production and applying software engineering discipline through modular design, version control, automated testing, documentation, CI/CD, release management, and production support.
- **Monitor and continuously improve live solutions** across reliability, response quality, performance, usability, security, and cost; troubleshoot production issues while applying responsible AI, security, privacy, and data-handling requirements throughout solution design, development, and operation.
- **Demonstrate independence, self-learning, accountability, and end-to-end ownership**, proactively communicating and managing technical risks, dependencies, blockers, and issues.

Essential Requirements:

- Bachelor's degree in Computer Science, Artificial Intelligence, Machine Learning, Data Science, Engineering, or a related technical field, or equivalent practical experience.
- 8-12 years of overall software engineering, AI engineering, machine learning engineering,

data engineering, or related technology experience, including significant hands-on experience delivering production applications.
- Strong programming capability in Python and practical experience using Jupyter Notebook for AI experimentation and development.
- Professional software development experience in Python and at least one additional language such as JavaScript/TypeScript or Java.
- Hands-on working experience with AWS and developing or deploying cloud-based applications.
- Proven hands-on experience designing and implementing RAG applications, including embeddings, chunking, retrieval, reranking, metadata filtering, context management, grounding, and hallucination mitigation.
- Hands-on experience with LangChain and LangGraph for LLM applications, agentic workflows, or multi-step orchestration.
- Hands-on experience with AI/LLM evaluation and observability using LangSmith, Ragas, or equivalent tools and methods.
- Practical working experience with Model Context Protocol (MCP), including building or consuming MCP-enabled integrations.
- Strong understanding of REST APIs, API-driven architecture, API Gateways, authentication/authorization, API security, rate limiting, and enterprise integration patterns.
- Experience building user-facing web applications with React.js, Next.js, or similar modern front-end frameworks.
- Experience taking AI capabilities from prototype through production, with testing, monitoring, documentation, and operational support.

Desirable Requirements:

- Hands-on experience with Amazon Bedrock, including foundation model APIs, Knowledge Bases, Agents, Guardrails, or model evaluation capabilities.
- Working knowledge of Databricks, Snowflake, Amazon S3, vector databases, hybrid search, and enterprise data platforms.
- Knowledge of multi-agent patterns, AI security, responsible AI, prompt and context engineering, and LLM observability.
- Relevant AI, Generative AI, cloud, or software engineering certifications. Examples include AWS Certified Machine Learning Engineer - Associate, AWS Certified AI Practitioner, AWS Certified Developer - Associate, Databricks Certified Generative AI Engineer Associate, or relevant Microsoft Azure AI credentials. Anthropic/Claude developer training or credentials are also beneficial where applicable.
- Experience working with cross-functional stakeholders, including product managers, architects, data engineers, UX teams, IT, security, and business users in Agile delivery environments.
- Exposure to Corporate Affairs, healthcare, pharmaceutical, communications, intelligence, or content-centric business processes is preferred but not essential.

Why Novartis: Helping people with disease and their families takes more than innovative science. It

takes a community of smart, passionate people like you. Collaborating, supporting and inspiring each other. Combining to achieve breakthroughs that change patients ' lives. Ready to create a brighter future together? https://www.novartis.com/about/strategy/people-and-culture

Benefits and Rewards: Learn about all the ways we ' ll help you thrive personally and professionally. Read our handbook (PDF 30 MB)

部门
Corporate Affairs

Business Unit
Communications

地点
India

站点
Hyderabad (Office)

Company / Legal Entity
IN10 (FCRS = IN010) Novartis Healthcare Private Limited

Functional Area
Data and Digital

Job Type
Full time

Employment Type
Regular

Shift Work
No

var kPlayer = KalturaPlayer55802022 || KalturaPlayer; var config = { targetId: "kalturaplayer6b05afa7d9244411620451", provider: { widgetId: "10m7rm1pm", partnerId: "2076321", uiConfId: "55802022" }, playback: { autoplay: false, autopause: false, muted: false, loop: false }, sources: { options: {}, startTime: 0 }, disableUserCache: "true", plugins: {}, sources: { options: {}, startTime: 0 }, ui: { showCCButton: false, settings: { showQualityMenu: true, showSpeedMenu: false }, components: { fullscreen: { disableDoubleClick: false } }, uiComponents: [{ presets: ['Playback', 'Live'], area: 'BottomBarRightControls', replaceComponent: 'Fullscreen', get: kPlayer.ui.components.Remove }] } }; // Check and add plugins only if they exist if (kPlayer.plugins["download"]) { config.plugins.download = { disable: true }; } if (kPlayer.plugins["transcript"]) { config.plugins["playkit-js-transcript"] = { position: "right", // Default: bottom;('left', 'right', 'top', 'bottom') to enable transcript. expandMode: "over", // Default: alongside;('alongside', 'hidden', 'over') expandOnFirstPlay: false, showTime: true, downloadDisabled: false, printDisabled: false, disable: true }; } if (kPlayer.plugins["preventSeek"]) { config.plugins.preventSeek = { preventSeekForward: false, preventSeek: false }; } config.plugins.floating = { disable: true }; if (kPlayer.plugins["navigation"]) { config.plugins.navigation = { position: "right", expandMode: "over", expandOnFirstPlay: false, visible: false }; } if (kPlayer.plugins["hotspots"]) { config.plugins['playkit-js-hotspots'] = { disable: true }; } if (kPlayer.plugins["moderation"]) { config.plugins['playkit-js-moderation'] = { disable: true }; } if (kPlayer.plugins["info"]) { config.plugins['playkit-js-info'] = { disable: true }; } if (kPlayer.plugins["share"]) { config.plugins.share = { disable: true }; } config.ui.uiComponents = []; if (kPlayer.plugins["googleAnalytics"]) { config.plugins.googleTagManager = {}; config.plugins.googleTagManager.customEventsTracking = {}; config.plugins.googleTagManager.containerId = 'GTM-57RJQ5'; config.plugins.googleTagManager.customEventsTracking.custom = []; config.plugins.googleTagManager.customEventsTracking = { preset: { coreEvents: true, UIEvents: false, playlistEvents: false, castEvents: false } }; }

// Ensure the global player registry array always exists, regardless of embed type. window.kalturaPlayerVideos = window.kalturaPlayerVideos || []; try { var thumbEmbedPromise = thumbnailEmbed({config, mediaInfo: {entryId: "1dgfvmafo"}}); // thumbnailEmbed() returns a Promise that resolves with the player instance // when the user clicks the thumbnail. Use .then() to capture the player directly. thumbEmbedPromise .then(function(player) { window.kalturaPlayerVideos.push(player); // Notify kalturadatalayer.js that a new player is ready so it can // attach custom event listeners immediately, regardless of when // the user clicked the thumbnail relative to page load. document.dispatchEvent(new CustomEvent('kalturaPlayerReady', { detail: { player: player } })); }) .catch(function(error) { console.error(error); }); } catch (e) { console.error(e.message) }

Accessibility and accommodation

Novartis is committed to working with and providing reasonable accommodation to individuals with disabilities. If, because of a medical condition or disability, you need a reasonable accommodation for any part of the recruitment process, or in order to perform the essential functions of a position, please send an e-mail to diversityandincl.india@novartis.com and let us know the nature of your request and your contact information. Please include the job requisition number in your message.

Job ID
REQ-10093342

AI Engineer

Apply to Job



Job ID
REQ-10093342

AI Engineer

Apply to Job

Source URL:
https://www.novartis.com.cn/careers/career-search/job/details/req-10093342-ai-engineer

List of links present in page

1. https://www.novartis.com/about/strategy/people-and-culture
2. https://www.novartis.com/sites/novartiscom/files/novartis-life-handbook.pdf
3. mailto:diversityandincl.india@novartis.com
4. https://novartis.wd3.myworkdayjobs.com/en-US/NovartisCareers/job/Hyderabad-Office/AI-EngineerREQ-10093342
5. https://novartis.wd3.myworkdayjobs.com/en-US/NovartisCareers/job/Hyderabad-Office/AI-EngineerREQ-10093342