QUESS
WINNING TOGETHER

Decoding India's AI Talent Landscape

About the Report

As AI moves from experimentation to enterprise-wide deployment, the nature of talent demand is evolving. This report captures that transformation, offering a panoramic view of the Indian AI and data talent landscape between March 2024 and March 2025.

It chronicles how organizations are recalibrating their talent strategies in response to GenAI's operationalization, the modernization of data ecosystems, and increasing pressure to innovate at scale. From the rise of GenAI Engineers and ML Ops Specialists to the consolidation of BI roles and the reinvention of traditional data science, the report unpacks the forces redefining job roles and skill expectations.

With sharp insights into hiring hotspots, skill scarcity, and wage differentials, the report helps business leaders, talent strategists, and policymakers decode where momentum lies — and what it will take to compete in the next wave of AI-driven growth.

Key Findings at a Glance – India's AI Talent Landscape

India's AI and Data talent market is rapidly changing, led by the rise of GenAI, modern data systems, and AI-focused product design. Job postings in AI/Data roles grew by 38–45% year-over-year, with the highest demand in Q3 FY25 (Oct–Dec 2024). GenAI-specific roles grew the fastest at 178%, signalling a major shift in role types and required skills.

Key Role Trends

- GenAI roles like Prompt Engineers and LLM Ops Specialists are emerging, replacing traditional AI/ML paths.
- Data Engineering grew by 61%, becoming central to GenAI and cloud platforms, with rising demand for tools like Airflow, Dagster, and vector databases.
- AI/ML roles grew by 42–43%, especially in industries like BFSI, Healthcare, and Retail.
- New roles like MLOps, AI Governance, and AI Product Management are on the rise.
- Cloud Engineers now focus on AI deployment, GPU scaling, and tools like Vertex AI and Bedrock.
- Traditional BI hiring fell by 14–16%, but there's more demand for AI-powered analytics and explainable dashboards.

Talent Supply Gaps – Upskilling & Internal Mobility Key

- Big Data engineers are moving into GenAI roles, up 26% YoY, showing strong reskilling momentum.
- Scarcity is highest for:
  - GenAI Engineers – only 1 candidate per 10 jobs
  - MLOps Specialists
  - Cloud Data Architects
  - AI Governance Experts

Pay and Location Insights

- GenAI roles pay 15–20% more than standard ML roles in metros.
- Tier-1 cities like Bengaluru, Hyderabad, and Pune lead in demand.
- Tier-2 cities offer lower costs but have limited talent availability; Kochi, Ahmedabad, and Coimbatore are gaining traction.

Employer Trends

- Global Capability Centres (GCCs) are driving hiring, shifting focus from cost-saving to innovation, personalization, and AI-led transformation.

In-Demand Skills for FY2026

- GenAI (Prompting, RAG, LLMOps)
- Advanced Data Engineering (Orchestration, Vector DBs)
- MLOps & AI Lifecycle
- Cloud AI (GPU scaling, cost-aware deployment)
- AI Governance & Risk
- AI Product Management
- AI-Augmented Analytics

Table of Content

Kapil Joshi
CEO, Quess IT Staffing

From The Leadership Desk

We are at a pivotal moment. **Artificial Intelligence is no longer a future possibility, it is a present reality.** And Generative AI is leading the charge, moving from experimental use cases to enterprise-wide transformation. It is reshaping how businesses operate, how teams are structured, and what capabilities organizations need to stay competitive.

Over the past year, the shift has been undeniable. Organizations have moved from curiosity to conviction. AI is now embedded across functions — from customer experience to fraud analytics to operational decision-making. And as adoption scales, the **talent landscape is transforming just as fast.**

Between March 2024 and March 2025, demand for AI and Data talent in India grew by an estimated **38% to 45%.** But the growth has been uneven. **GenAI-specific roles** — Prompt Engineers, GenAI Researchers, LLMOps Specialists — surged by over **170%**, reflecting the urgency to productionize AI capabilities. **Data Engineering** has become the foundation of scalable AI, while traditional BI roles are contracting, mirroring automation and toolchain consolidation.

It's not just job titles that are shifting — it's expectations. Engineers and analysts are now required to build LLM-ready data pipelines, work across orchestration frameworks, and deploy AI systems in real time. The classic "Data Scientist to ML Engineer" pathway has fragmented into specialized, cross-functional tracks.

As roles evolve, **scarcity has become acute.** In GenAI Engineering, there's roughly one qualified candidate for every ten open roles. MLOps, AI Infrastructure, and Governance roles show similar pressure. This isn't just a hiring challenge — it's a strategic one.

Yet there's momentum. More professionals are transitioning internally, Big Data Engineers into GenAI pipelines, Cloud Engineers into AI-native platforms. This shift from hire-to-fill to grow-to-fill is becoming a core capability.
Geographically, **Tier-1 cities** like Bengaluru, NCR, and Hyderabad remain dominant, but **Tier-2 hotspots** like Kochi, Coimbatore, and Ahmedabad now account for nearly **70% of Tier-2 demand**. These markets offer scale and cost efficiency, though senior AI talent remains limited.

Perhaps the most significant shift is in **Global Capability Centres (GCCs)**. They are no longer just operational support arms, they are driving AI innovation. From scaling GenAI platforms to building governance frameworks, GCCs are anchoring enterprise-wide AI readiness.

In all of this, one truth holds: **it is people who power progress.** This report is not just a snapshot — it's a compass to help you navigate what's ahead.

Introduction

India's AI and data talent ecosystem is experiencing a structural shift. Over the past 12 months, what began as experimentation with Generative AI (GenAI) has evolved into full-scale deployment across industries. As enterprises move from prototypes to production, the talent landscape is adapting — not only in volume of demand but in the depth and specificity of skills required.

Between March 2024 and March 2025, hiring demand for AI and data professionals grew by an estimated **38% to 45%**. However, this growth has been uneven. Some job families, particularly those tied to GenAI architecture, orchestration, and deployment, are accelerating rapidly. Others, like traditional Business Intelligence (BI) roles, are contracting or being absorbed into newer platform workflows.

At the same time, the nature of roles is evolving. There is a clear movement away from generalist AI profiles toward highly specialized functions, in areas like prompt engineering, LLM operations, retrieval-augmented generation (RAG), and model evaluation. These changes demand not only new hiring strategies but also significant internal reskilling and cross-functional team design.

This report analyses these shifts through a multi-dimensional lens: demand growth, talent supply constraints, skill evolution, geographic spread, and employer segmentation. It aims to help talent leaders, CHROs, and strategy teams make sense of where momentum lies — and what the emerging shape of the AI workforce looks like.

Highlights from the year in review

- **+178% YoY growth** in GenAI-specific roles (Prompt Engineers, GenAI Scientists, LLMOps)
- **+61% YoY growth** in Data Engineering roles, now critical to scalable AI infrastructure
- **14–16% decline** in BI hiring as automation and toolchain consolidation reshape the space
- **High to severe scarcity** for GenAI Engineering and MLOps profiles, with talent availability falling below 1:10 in top cities
- **70% of Tier-2 demand** concentrated in Kochi, Coimbatore, and Ahmedabad — indicating the start of regional diversification
- **GCCs driving enterprise-wide hiring** across AI/ML and data roles, reflecting a pivot from execution to innovation ownership

The following sections examine these patterns in depth — role by role, skill by skill, and city by city. This report is not just a reflection of what's changed — it is a blueprint for how to build what's next.

Market Overview

The AI and data talent landscape is undergoing a significant transformation. Job families are evolving in different directions—some are expanding, others consolidating, and some are being fundamentally redefined. Overall, job growth has increased by 38%-45% year-over-year.

Highlights from the year in review

- Generative AI roles experienced significant growth, increasing by 1.8x.
- Data Engineering remains fundamental, serving as the backbone for AI and GenAI projects.
- Newer categories of roles are emerging, such as AI Product Managers and AI Governance Leaders.
- Highly in-demand emerging profiles include MLOps, AI Lifecycle Specialists, and GenAI Architects.
- A significant percentage of AI-related postings in India now require skills in prompt engineering, model evaluation, or RLHF tuning (16-20% in early 2025).
- Data engineering roles increasingly require expertise in orchestration tools (Airflow, Dagster) and vector databases (Pinecone, Weaviate, Qdrant), indicating convergence with GenAI infrastructure needs (18-22% explicitly mention these).
- Lateral movement of Big Data engineers into GenAI infrastructure roles has expanded by 26% YoY, underscoring the need for rapid re-skilling.
- Cloud engineering job descriptions increasingly include AI deployment responsibilities like GPU cluster management and scalable LLM operations (more than 16%).
- While overall BI hiring volumes have declined (14-16%), there is rising demand for AI-augmented analytics engineers capable of embedding explainable AI within traditional reporting.

Macro Talent Trends

The ecosystem is driven by GenAI operationalization, modern data architecture, and AI-native design — leading to divergent job family trends. Data Engineering and AI/ML roles are expanding, while BI & Warehousing consolidates.

Category Insights

- Data Engineering: Plays a key role in building the systems needed to support AI, helping manage data flow and storage.
- AI/ML (including Generative AI): The field is expanding with new types of jobs, including those focused on advanced AI models and research.
- Cloud: Efforts are focused on running AI tools efficiently in the cloud, keeping costs in check, and ensuring reliable performance across platforms.
- Business Intelligence (BI) & Data Warehousing: These areas are becoming more streamlined and automated, leading to fewer traditional BI roles.
- Data & Analytics: Moving toward more specialized roles that focus on using data to predict outcomes and support real-time decision-making within specific business areas.

Job Role Evolution Matrix

As India's AI ecosystem matures, job roles are undergoing visible structural shifts. Traditional linear career paths — such as Data Scientist to ML Engineer — are fragmenting into specialized tracks aligned to Generative AI (GenAI), infrastructure orchestration, and model lifecycle management. At the same time, employer expectations are rising. Roles are becoming more cross-functional, tool-agnostic, and production-focused.

Between March 2024 and March 2025, three categories of role movement became clear:

- Rapid expansion of GenAI-native roles, driven by productionization of LLMs
- Repositioning of infrastructure and orchestration roles, particularly around data and cloud
- Consolidation of traditional roles, such as BI Developers, amid platform automation

These shifts are not theoretical. These changes are already visible in job postings, salary bands, and tool expectations across Tier-1 and Tier-2 cities — confirming that AI hiring is now shaped by specialization and deployment readiness.

Traditional Role	Evolved / Emerging Role(s)	Evolution Driver	Growth / Decline Trend
Data Scientist	- GenAI Research Scientist - Model Evaluation Lead - RLHF Specialist	Performance tuning, safety, model alignment	Stable; specialization emerging
ML Engineer	- LLMOps Engineer - RAG Pipeline Engineer - Inference Optimization Lead	Model deployment complexity, retrieval workflows	+42% to 43% YoY
Data Engineer	- Vector Data Engineer - Orchestration Specialist (Airflow, Dagster) - LLM-Ready Pipeline Engineer	Shift to real-time, token-efficient pipelines	+61% YoY
BI Developer / Analyst	- Augmented Analytics Engineer - Data Product Designer	Automation of reporting, dashboard consolidation	–14% to –16% YoY
Cloud / DevOps Engineer	- AI Infrastructure Engineer - GPU Cluster Manager - AI Cost Ops Lead	AI-native cloud orchestration, GenAI scaling on cloud	Growing; 16% of postings now AI-tagged
Big Data Engineer (Transition)	- GenAI Infrastructure Engineer - Feature Store Developer	Adjacent skill mobility, data readiness for GenAI	+26% lateral movement
AI Product Manager	- GenAI Product Owner - AI Platform Strategist	Business alignment, full-stack AI delivery	Emerging; high GCC demand
Compliance Analyst / IT Auditor	- Responsible AI Consultant - AI Governance Lead	Explainability, risk, ethics, regulatory alignment	Emerging; flagged for 2025 growth

Key Insights

- GenAI specialization is redefining core AI job tracks. Roles such as Prompt Engineers and GenAI Researchers now account for over 10% of new AI job postings, a significant jump from under 2% in 2023.
- Data Engineering has become the backbone of GenAI infrastructure. Today, 18–22% of job descriptions for data engineers explicitly reference orchestration tools (Airflow, Dagster) alongside LLM-specific capabilities.
- Cloud engineers are expected to own AI infrastructure. Over 16% of cloud job postings now include responsibilities for AI deployment, GPU orchestration, or LLM stack integration — up from ~8% a year prior.
- Internal transitions are rising. Lateral movement from Big Data to GenAI Infrastructure roles has grown by 26% YoY, often enabled through targeted upskilling in vector databases, embeddings, and orchestration layers.
- These trends signal a fundamental shift: AI hiring is no longer driven by broad capability labels — but by deep specialization aligned to deployment readiness, tool fluency, and lifecycle ownership.

The next section explores how these evolved roles translate into demand patterns across skill clusters, and where talent supply constraints are emerging most sharply.

Demand Analysis by Skill Cluster

India's AI and data talent market is no longer defined by headline growth alone — it is now shaped by where growth is concentrated, how roles are evolving, and what portion of total hiring each skill cluster represents.

Between March 2024 and March 2025, the market saw significant divergence across skill clusters. GenAI Engineering, while starting from a lower base, recorded the fastest growth. Meanwhile, Data Engineering and AI/ML Engineering continued to anchor the bulk of hiring volumes. In contrast, Business Intelligence and Visualization roles saw year-on-year declines, reflecting the ongoing automation and consolidation of legacy analytics functions.

This section summarizes the movement across high-growth, stable, emerging, and declining skill areas.

Skill Cluster Demand Trends – Growth and Share of Total Demand

Skill Cluster	Sample Roles	YoY Demand Growth	Share of Total Demand (FY24)	Trend
GenAI Engineering	Prompt Engineer, LLMOps, GenAI Scientist	+178%	10–12%	Explosive growth, building volume
Data Engineering	Vector Specialist, Airflow/Dagster Engineer	+61%	28–30%	Largest cluster, critical to LLM readiness
AI/ML Engineering	Model Evaluator, RAG Engineer, ML Engineer	+42–43%	22–24%	High demand, sector-wide adoption
MLOps & AI Infra	GPU Ops, AI Cost Optimizer, AI DevOps	~+20%	8–10%	Stable, emerging as core capability
AI Product Management	GenAI PM, Use Case Strategist	Emerging	2–3%	Early stage, led by GCCs
AI Governance & Risk	Responsible AI Lead, Model Audit Consultant	Emerging	2–3%	Niche but growing, especially in BFSI
BI & Visualization	BI Developer, Dashboard Specialist	–14 to –16%	15–17%	Declining, increasingly absorbed by AI teams

GenAI Engineering shows steep growth, but Data and ML Engineering still anchor over 50% of AI/Data hiring — underpinning most infrastructure pipelines. Meanwhile, MLOps is stabilizing, and BI roles continue their downward trend due to automation.

BI and Visualization roles, once core to enterprise analytics teams, are contracting as tools consolidate and automation platforms absorb legacy reporting workflows.

Year-on-Year Growth in Hiring Demand by Skill Cluster

YoY Growth (%)

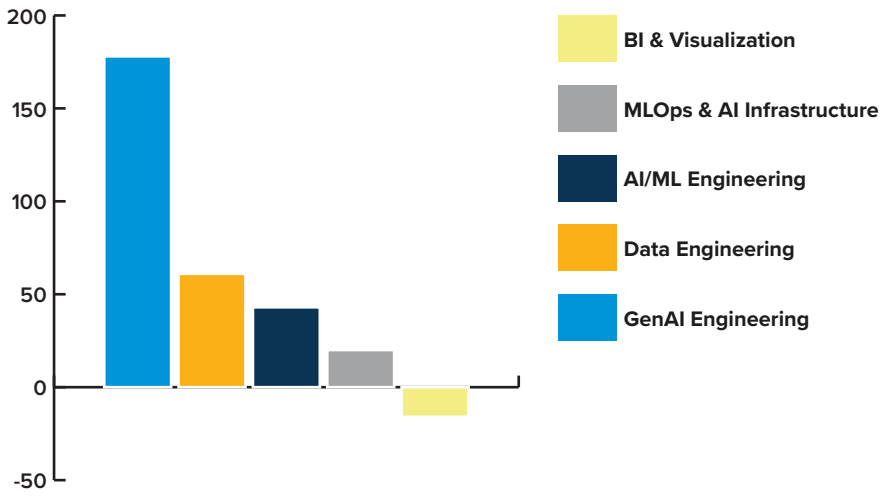

GenAI roles led hiring momentum, while Data Engineering and ML Engineering continued to anchor volume.

Share of Total AI/Data Demand by Cluster

Share of Total Demand (%)

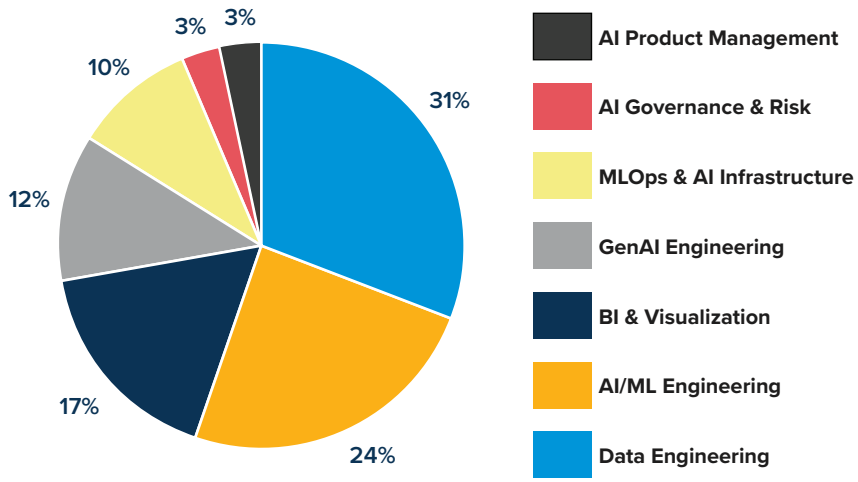

Despite rapid growth, GenAI roles remain a smaller slice of overall hiring — while Data and ML Engineering dominate volume.

India's AI talent demand is evolving along two dimensions: vertical depth (growth rate) and horizontal maturity (share of total demand). Organizations are no longer hiring for broad skill clusters — they are building teams aligned to stack architecture, model lifecycle, and infrastructure readiness.

In the next section, we analyze how these shifts map to employer types — and how different organizations are shaping their AI hiring strategies.

Tools, Skills, and Signals

As AI roles become more specialized, hiring is shifting from general experience to **stack readiness**—a blend of tools, applied skills, and the ability to deliver in production. Between March 2024 and March 2025, job descriptions across GenAI, data, cloud, and analytics began explicitly listing the **platforms** and **practices** required. These tool-skill pairings reflect not just role expectations, but also where a company stands in its AI journey—from experimentation to scaled deployment.

The table below summarizes high-demand role clusters, their key technical skills, and the most frequently cited platforms in hiring.

Role Cluster to Skills and Tools Mapping

Role Cluster	Key Skills / Practices	Common Tools / Platforms
GenAI Engineering	Prompt engineering, model evaluation, RAG workflows, RLHF tuning	OpenAI API, LangChain, HuggingFace, Gradio, RAG, Transformers
Data Engineering	Pipeline orchestration, vector ingestion, token optimization	Apache Airflow, Dagster, dbt, Snowflake, Pinecone, Weaviate
AI/ML Engineering	Feature engineering, model experimentation, lifecycle tracking	Scikit-learn, TensorFlow, PyTorch, MLflow, XGBoost, DVC
MLOps & Infra	Model deployment, observability, reproducibility, scalability	Docker, Kubernetes, Ray, FastAPI, Weights & Biases
Cloud AI Engineering	GPU provisioning, cost optimization, scalable LLM hosting	Vertex AI, SageMaker, Azure OpenAI, GCP/AWS, Terraform
BI & Analytics	Conversational UX, LLM-augmented dashboards	Power BI, Tableau, Looker, ChatGPT API, LangChain UI
AI Governance & Risk	Explainability tuning, bias detection, fairness auditing	SHAP, LIME, Fairlearn, IBM Watson OpenScale, EBM

Key Insights

- 18–22% of Data Engineering job postings referenced orchestration or vector DB tools (Airflow, Dagster, Pinecone) — showing infrastructure maturity for GenAI deployment.
- 16% of Cloud Engineering roles now include AI-specific infrastructure expectations such as GPU orchestration and cost-efficient scaling of LLMs.
- Skills such as prompt engineering and retrieval-augmented generation (RAG) are now embedded in both GenAI and BI roles, showing a convergence of user-facing and infrastructure-facing capabilities.
- Governance tools like SHAP, LIME, and Fairlearn are becoming more visible in BFSI and healthcare JDs — often attached to risk, compliance, or platform governance roles.

Strategic Takeaway

AI hiring is now defined by stack specificity, talent is assessed not just on theory, but on hands-on fluency with modern AI toolchains. For organizations scaling GenAI, this mapping provides a clear blueprint for role design, hiring, and reskilling.

The next section explores which employer types are fueling demand for these tool-aligned roles, and how their priorities are shaping volume and capability expectations.

Employer Landscape

India's AI hiring momentum is increasingly shaped by a blend of employer types — each bringing different levels of digital maturity, GenAI ambition, and platform ownership. While tech firms and startups continue to innovate at the edge, it is Global Capability Centres (GCCs) that are leading high-complexity, GenAI-enabled hiring in India.

Between March 2024 and March 2025, GCCs, system integrators, and product firms together accounted for over 70% of all AI and data-related hiring in the country. However, demand is no longer limited to digital-first organizations. Sectors like BFSI, retail, healthcare, and manufacturing — especially through GCCs and captives — are scaling their GenAI capabilities, creating demand for specialized infrastructure, orchestration, and governance roles.

AI & Data Hiring Share by Employer Type

Employer Type	Share of Total AI/Data Hiring	Hiring Focus
System Integrators (IT Services)	38.0%	Deployment scale, MLOps integration, cloud modernization
Global Capability Centres (GCCs)	22.5%	GenAI teams, platform engineering, AI lifecycle roles
Product Development Firms	12.0%	Feature-led innovation, inference ops, GenAI-driven personalization
Consulting & Analytics Firms	10.5%	Advisory, compliance AI, explainability and governance hiring
Startups / DeepTech Firms	6.5%	Prototyping, LangChain development, GenAI app experimentation
Captive Enterprise AI Labs	3.5%	Domain-specific R&D, LLM tuning, internal transformation pilots
Others (Non-Tech Enterprises)	7.0%	Early-stage analytics and AI experimentation

System integrators dominate in terms of volume, while GCCs and product firms increasingly anchor high-value, high-complexity roles aligned to GenAI production.

Share of Total AI/Data Hiring

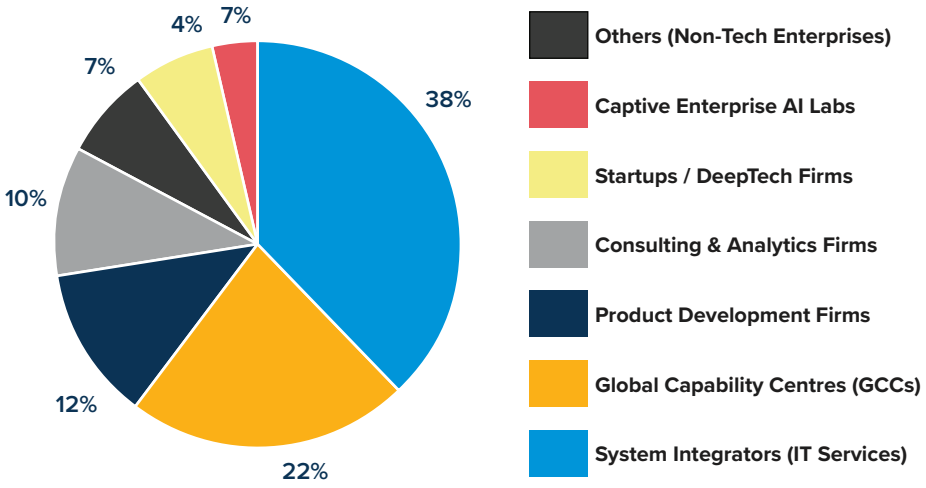

GCCs: The Strategic AI Engine

Though accounting for 22.5% of overall AI/data demand, GCCs are playing an outsized role in shaping India's GenAI workforce. With the shift from execution to **platform ownership,** GCCs are rapidly scaling specialized teams in LLMOps, PromptOps, model lifecycle engineering, and cloud-native GenAI delivery.

A sectoral view of GCC AI hiring highlights the breadth and depth of their transformation efforts:

AI & Data Jobs Distribution Across Indian GCC Sectors

GCC Sector	Indicative % Share	Hiring Focus / AI Maturity
Technology, Hardware & Internet	24%	AI/ML CoEs, platform-level ownership, LLM production ops
BFSI	22%	Data Governance, Risk Analytics, GenAI for compliance and automation
Retail & eCommerce	11%	Personalization engines, chat-based UX, real-time decisioning
Telecom & Media	11%	Digital CX, GenAI copilots, NLP interfaces for customer engagement
Manufacturing, Auto, Energy	10%	Predictive ops, AI for supply chains, ESG & industrial AI pilots
Healthcare & Pharma	7%	Diagnostic AI, GenAI for trial design, regulatory compliance
Professional Services	7%	Big 4 analytics centers, AI consulting hubs
Travel, Hospitality & Logistics	3%	Route optimization, pricing AI, automation of support ops
Defence & Aerospace	2%	Specialized, high-security GenAI teams
Construction & Infrastructure	1.5%	Smart city AI, predictive maintenance, digital twin integration
Real Estate & PropTech	1.5%	LLM-assisted customer experience, lead prediction

Indicative % Share

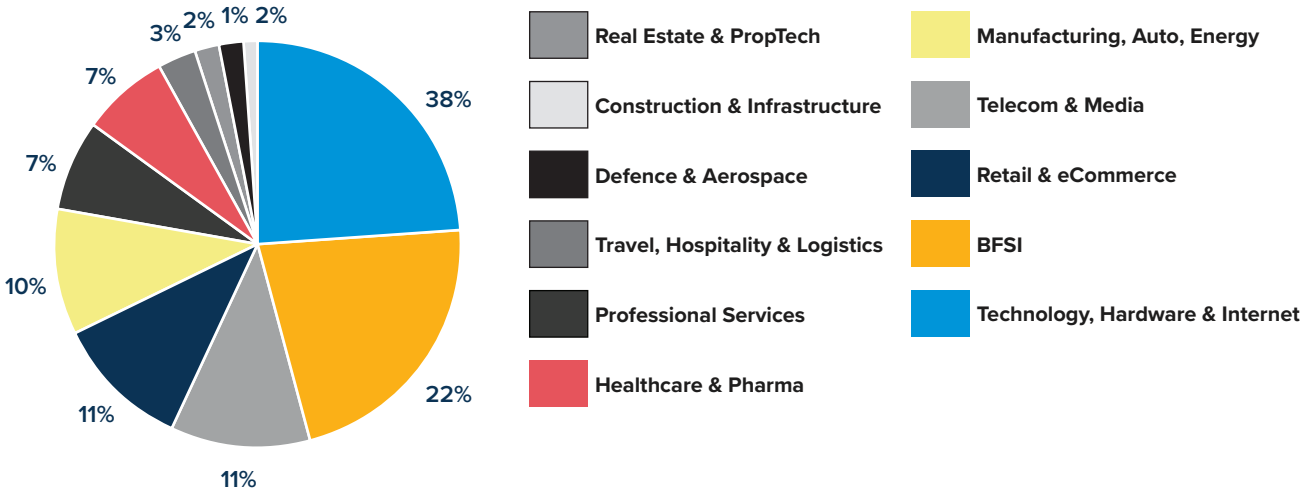

Across these sectors, GCCs are building **AI-native transformation units** — many structured as internal platforms serving global business units. Their scale, access to India's deep talent pool, and cloud maturity are accelerating the institutionalization of GenAI adoption.

Sectoral Employer Priorities

- Product firms (12%) are embedding GenAI across user journeys, chat interfaces, and platform intelligence. Roles like Prompt Engineers, LLMOps Specialists, and GenAI UX Leads are common.
- System Integrators (38%) anchor volume. Hiring focuses on MLOps Leads, Cloud Data Engineers, and enterprise-ready GenAI integrations for clients.
- Consulting firms (10.5%) are shaping AI strategy and governance frameworks — particularly in BFSI, pharma, and insurance — with roles in model audit, explainability, and AI regulation.
- Startups and DeepTech (6.5%) serve as early innovation labs, pushing boundaries with HuggingFace fine-tuning, RLHF experimentation, and LangChain prototyping.
- Captives and enterprise labs (3.5%) are building domain-aligned use cases — especially in risk modeling, HR automation, and knowledge retrieval.

Strategic Takeaway

The Indian AI employer landscape is undergoing a structural shift — from project staffing to platform engineering, and from PoCs to production. GCCs, product firms, and strategic consultancies are leading the charge, while system integrators scale deployment and service breadth. The nature of hiring varies not just by industry, but by organizational intent and digital maturity.

In the next section, we examine how talent is responding, through lateral movement, cross-skilling, and the reshaping of internal mobility pathways in India's AI workforce.

Talent Movement & Reskilling

As GenAI reshapes the AI talent landscape, organizations across India are no longer relying solely on external hiring to close capability gaps. Instead, they are increasingly turning inward — redeploying, cross-skilling, and fast-tracking internal talent into newly emerging roles.

Between March 2024 and March 2025, the talent market saw a marked increase in lateral movement across adjacent AI job families, coupled with the rise of structured reskilling programs aligned to evolving tech stacks. These shifts reflect not only talent scarcity but a growing emphasis on workforce adaptability.

Talent Flow – Internal Movement Trends

While Section 5 outlined how roles are structurally evolving, talent is actively responding. Internal mobility data shows clear patterns of lateral movement, particularly where foundational technical skills are transferable and GenAI readiness can be developed through focused training.

From	To	YoY Lateral Movement	Why It's Happening
Big Data Engineer	GenAI Infrastructure Engineer	+26%	Compatible orchestration and data architecture skills
BI Developer	AI-Augmented Analyst (LangChain, ChatGPT API)	+18%	BI roles are consolidating; GenAI dashboards expanding
Cloud/DevOps Engineer	Cloud AI Engineer / LLM Deployment Specialist	+14%	Infrastructure maturity, rise of GPU orchestration
Data Scientist	Model Evaluator / Prompt Engineering	+13%	Shift from model building to model alignment and tuning
ML Engineer	LLMOps / RAG Pipeline Architect	+11%	Retrieval workflows and scalable deployment expectations

These lateral transitions are less about role rebranding and more about stack adaptation — as emerging clusters like GenAI Infrastructure, LLMOps, and Prompt Engineering demand new tool fluency but build on adjacent capabilities.

Reskilling Velocity and Capability Tracks

Employers, particularly GCCs and product firms, are investing in modular reskilling programs designed to convert adjacent-role talent into GenAI-ready contributors in 8–12 weeks. These are not generic MOOCs, but stack-specific capability tracks tailored to internal workforce profiles.

Common pathways include:

- Prompt Engineering: Positioned as a modular add-on for analysts, content QA teams, and junior engineers
- LangChain & RAG Workflows: Targeted at BI developers and data engineers moving into conversational or search-augmented pipelines
- Vector Database & Embedding Fundamentals: Now a standard part of advanced data engineering upskilling tracks
- LLMOps Toolchains: Kubernetes for LLM orchestration, ML observability, inference cost tracking
- AI Governance & Explainability: Cross-trained from compliance, audit, and risk analytics backgrounds

In several GCCs and large services firms, cross-skilling pilots are now being tied directly to internal job rotations, often backed by performance-linked incentives or career path revisions.

Strategic Implications

These shifts have tangible implications for workforce strategy:

- Lateral redeployment is now core to AI hiring strategy, particularly where GenAI-ready external talent is scarce
- Reskilling velocity is a competitive differentiator — organizations able to retrain in <3 months are outperforming slower talent ecosystems
- Internal mobility is reshaping employer branding — fast-growing AI teams increasingly promote upskilling opportunities as part of EVP

Rather than shrink, the AI workforce is remapping itself, from legacy role families to new capability clusters defined by stack ownership, tooling fluency, and delivery readiness.

In the next section, we explore how these reskilling and hiring patterns are reflected geographically, with shifts in demand intensity across Tier-1 and Tier-2 cities, and new regional hotspots emerging in response to infrastructure, wage arbitrage, and GenAI workforce maturity.

Location-Based Talent Insights

India's AI and GenAI hiring landscape is becoming geographically stratified. While Tier-1 metros remain dominant in both hiring volume and capability depth, Tier-2 cities are emerging as secondary hubs for analytics delivery, cloud engineering, and early-stage experimentation. These geographic distinctions now play a critical role in talent planning, cost optimization, and capability building.

Tier-1 AI Hiring Volume and Talent Sufficiency

The table below outlines the distribution of AI and GenAI talent and job postings across India's major Tier-1 cities, along with their demand intensity and talent sufficiency indicators.

City	% of AI Talent	% of Job Posts	Hiring Intensity	Supply Sufficiency Index
Bengaluru	29%	31%	Very High	3.9
NCR	18%	12%	High	4.1
Hyderabad	12%	12%	High	4.0
Mumbai	8%	15%	Very High	3.13
Chennai	8%	8%	High	3.8
Pune	6%	8%	High	4.2
Kolkata	4%	3%	High	3.33
Others	15%	11%	Moderate	4.2

Interpretation: Bengaluru remains India's AI epicenter, accounting for nearly a third of both talent supply and hiring. However, it also presents one of the most competitive hiring environments. Mumbai shows a higher job demand relative to its available talent, particularly in BFSI GenAI applications. Chennai offers a balanced supply-demand environment, while Pune is seeing moderately elevated demand for its available pool.

Tier-2 Cities: Demand Distribution and Growth Signals

Tier-2 cities now account for approximately 14–16% of India's AI/ML hiring demand. Although GenAI specialization is still nascent, these cities offer cost-effective pathways for scaling mid-senior delivery teams, provided infrastructure and training investments are made.

City	Share of Tier-2 Demand (%)	Primary Drivers
Kochi	27.3%	Data Engineering, cloud CoEs, startup-led R&D
Ahmedabad	22.7%	BFSI analytics hubs, early GCC expansions
Coimbatore	18.2%	Early-stage AI product hiring, BI modernization
Jaipur	13.6%	Automation delivery, analytics support centers
Indore	9.1%	Cloud Data Engineering workflows, annotation ops
Chandigarh/Mohali	6.8%	Uptick in public sector and GovTech AI hiring
Others	2.3%	Scattered demand in cities like Bhopal, Nagpur, Visakhapatnam

Insight: Kochi, Ahmedabad, and Coimbatore make up nearly 70% of Tier-2 AI hiring. These locations are being increasingly leveraged for data engineering, BI modernization, and GenAI support workflows, especially by mid-sized GCCs and system integrators.

City-Level Role Specialization

Each metro is developing domain-specific AI capabilities driven by its sectoral footprint and infrastructure maturity.

City	Primary Specializations
Bengaluru	GenAI startups, cloud-native CoEs, full-stack AI platform teams
Hyderabad	Prompt Engineering, LLMOps, Cloud AI delivery (AWS, Azure, GCP)
NCR (Delhi)	AI Governance, RAG Engineering, BFSI fraud prevention
Mumbai	BFSI GenAI adoption, Risk AI, regulatory compliance, personalization AI
Chennai	AI in SaaS/PaaS, cloud transformation in BFSI and engineering services
Pune	Mid-office AI ops (DataOps, QA, ModelOps), digital transformation
Kolkata	BFSI data cleansing, ingestion, and entry-level analytics roles
Coimbatore / Ahmedabad / Jaipur	Captive AI/Data engineering for BFSI, manufacturing sectors

A Deep-Dive: Hiring Drivers by City

While hiring volumes differ widely, each major city in India is developing distinct functional niches, shaped by industry density, digital infrastructure, and proximity to innovation ecosystems.

City	AI/GenAI Hiring Drivers
Bengaluru	A major hub for GenAI startups, full-stack AI engineers, and platform teams. Cloud-native CoEs and SaaS product companies drive high LLM and Prompt Engineer demand.
NCR	AI Governance and BFSI use cases dominate. Roles in Risk AI, fraud detection, credit modelling, and regulatory analytics are in high demand.
Hyderabad	Strong Cloud CoE ecosystem (AWS, Azure, GCP). Demand is growing in LLMOps, Prompt Engineering, and fine-tuning. Mid- to large-sized IT firms are scaling GenAI ops.
Mumbai	BFSI drives adoption through GenAI in hyper-personalization, risk analytics, and predictive credit scoring. Explainability and regulatory compliance are critical.
Chennai	SaaS and engineering services are adopting AI platforms. BFSI GCCs are transforming data infra and embedded AI into core operations.
Pune	Insurance and fintech GCCs are expanding analytics engineering. Mid-office AI ops roles (DataOps, QA, ModelOps) are on the rise.
Kolkata	BFSI data cleansing and ingestion teams dominate. GenAI hiring is nascent. Entry-level data analytics, stewardship, and reporting roles prevail.
Coimbatore / Ahmedabad / Jaipur	Captive delivery hubs for BFSI and manufacturing. Cost-effective for mid-senior data engineering and BI modernization.

This granular mapping of hiring demand provides a practical lens for employers seeking to optimize location-strategy, whether for high-scarcity GenAI roles or scalable support functions.

10.4 Strategic Considerations

- Tier-1 cities remain the epicenter for AI infrastructure, platform engineering, and GenAI innovation
- Tier-2 hubs are expanding in delivery, BI modernization, and cloud data workflows — provided skill gaps are addressed
- Employers are increasingly adopting hub-and-spoke models, centralizing strategic talent and decentralizing repeatable operations
- Role-specific maturity by city must guide location planning, not just volume or cost metrics

India's Tier-1 cities are not just hotspots for AI hiring — they are hosting the strategic transformation of Global Capability Centres (GCCs). These centres are evolving rapidly from operational support units into GenAI-powered innovation engines. The table below outlines the key forces behind this shift:

Driver	Implication for GCCs
Strategic Shift from Cost Arbitrage to Innovation Ownership	GCCs are now tasked with leading innovation, product development, and digital transformation — not just support functions.
Enterprise-Wide Digital Transformation Alignment	GCCs are executing strategic AI and digital programs aligned with global roadmaps.
Availability of AI-Native Talent Pools	India's GenAI talent density allows rapid scaling from pilot to production.
Operational Efficiency and Agility Imperatives	GenAI enables automation of knowledge workflows and leaner operations.
Competitive Pressure among Peer GCCs	Early adopters secure larger remits and attract more strategic charters.
Rising Demand for Hyper-Personalized Experiences	GCCs are driving GenAI-based personalization in CX and EX journeys.
Risk Management and Compliance Automation	GenAI helps automate compliance, fairness, and governance frameworks.
Cloud-First Infrastructure Enabling GenAI Scalability	Mature cloud adoption allows faster scaling of LLMs and GenAI workflows.
Executive Mandate for AI-First Transformation	GCC KPIs now explicitly include GenAI delivery and maturity milestones.

Summary Takeaway

Location strategy is no longer about volume — it's about capability alignment. Tier-1 cities remain AI innovation engines, while Tier-2 hubs offer scale and cost agility for data delivery and support roles.
Next, we examine how these capability shifts and hiring dynamics translate into compensation structures and skill premiums, and what employers must plan for as demand continues to surge.

Compensation Trends & Premiums

As India's AI talent landscape matures, compensation structures are diverging sharply between traditional analytics roles and GenAI-specific functions. What was once a relatively flat salary band across titles is now a stack- and capability-driven market, where production-readiness and tool fluency drive value.

From March 2024 to March 2025, salary premiums emerged most prominently in roles tied to GenAI deployment, orchestration, and lifecycle ownership. These functions demand operational mastery of LLMs, vector databases, and cloud-native infrastructure — making them harder to source, and significantly more expensive to hire.

GenAI Roles Are Redefining Salary Norms

GenAI roles now command 15–20% higher compensation than traditional AI/ML roles, even at the same experience level. This is most visible in roles such as Prompt Engineer, GenAI Architect, and GenAI Research Scientist, where hands-on expertise with orchestration frameworks, vector databases, and model tuning is non-negotiable.

GenAI Role	Bengaluru (LPA)	Hyderabad (LPA)	Pune (LPA)	Chennai (LPA)
Prompt Engineer	37.0–52.0	34.0–47.0	31.0–43.0	30.0–41.0
GenAI Research Scientist	38.5–54.0	35.0–49.5	32.5–45.0	31.0–43.5
GenAI Solutions Architect	41.0–57.0	37.0–52.0	34.0–48.0	32.5–45.5
Principal AI Architect	42.0–58.0	38.0–53.0	34.5–49.0	33.0–46.0
Data Scientist (GenAI + NLP)	36.0–50.0	33.0–46.0	30.0–42.0	29.0–40.0

These bands reflect a combination of scarcity, stack dependency, and productization urgency — all of which elevate these roles beyond general ML profiles.

Traditional Roles Show Stability, With Select Uplift

Traditional AI, ML, and data platform roles continue to see stable demand and moderate uplift — particularly where these functions intersect with cloud engineering and lifecycle ops. However, they generally remain 10–20% below GenAI counterparts in terms of salary.

Traditional Role	Bengaluru (LPA)	Hyderabad (LPA)	Pune (LPA)	Chennai (LPA)
AI & Data Science Engineer	35.3–47.9	33.1–45.0	31.0–41.2	30.0–40.3
Machine Learning Engineer	34.0–46.0	31.5–43.5	29.0–40.0	28.5–38.5
AI/ML Research Scientist	36.0–51.0	33.0–47.0	30.5–44.0	29.5–42.0
AI Infrastructure Architect	40.0–56.0	36.0–51.0	33.5–47.0	32.0–45.0
Cloud Data Architect	36.0–48.5	33.0–45.5	30.0–42.5	28.5–40.0
Data Platform Engineer	34.5–46.5	31.0–43.0	29.0–40.0	28.0–38.0
DataOps Engineer	33.0–45.0	30.0–41.0	28.0–39.0	27.0–37.0
RPA Engineer (AI-Enabled)	31.0–42.5	28.5–39.5	26.0–36.5	25.0–34.5
BI Engineer (LLM-Integrated)	30.0–41.0	27.5–38.0	25.5–35.5	24.5–34.0
Data Analytics Specialist	32.5–43.5	30.0–40.0	28.0–37.0	26.5–35.5

These roles remain essential to GenAI enablement, but tend to anchor platform integration, rather than innovation. Pay reflects this distinction.

Strategic Observations

- Salary benchmarks are now capability-led, not title-led: Two engineers with similar experience but different stack exposure can show a 15–20% pay gap.
- Metro concentration remains strong for GenAI roles, particularly in Bengaluru and Hyderabad. Salary ceilings correspond with both tool demand and ecosystem maturity.
- Tier-2 deployment is feasible for traditional roles with proper skilling investment — but GenAI hiring in Tier-2 is still nascent.
- Pay inflation is slowing, but replacement costs for stack-specialists remain high — making retention strategy critical.

Summary Takeaway

India's AI salary landscape is stratifying around stack relevance and production readiness. GenAI roles now represent a new compensation tier, with distinct expectations and wage expectations. To hire competitively and scale responsibly, organizations must move from flat role bands to calibrated compensation frameworks that reflect capability, tool depth, and market scarcity — not just experience or designation.

AI Talent Scarcity

As demand for AI and GenAI roles accelerates across industries, a more persistent challenge is emerging — there simply aren't enough qualified professionals to fill the most critical roles. Compensation is rising, yes, but it's the deepening supply gap that defines FY25's hiring landscape.

This section combines role-specific scarcity ratings with a quarterly heatmap of talent shortfall, giving employers a clear view of where AI pipelines are most strained — and what capabilities require urgent investment.

Scarcity Index: Where the Market Is Breaking

India's talent shortage is most severe at the intersection of infrastructure, deployment, and governance. Roles that support GenAI execution — such as engineering, lifecycle management, and ethical oversight — face acute supply constraints.

Skill Cluster	Scarcity Level	Indicator
GenAI Engineers	Very High	~1 qualified candidate for every 10 open roles
MLOps Specialists	High	Persistent shortage of model deployment and lifecycle engineers
Cloud-Native Data Architects	Medium–High	Demand outpaces supply, especially in Tier-1 cities
AI Governance & Risk	Emerging Scarcity	Early-stage demand, but already limited availability
Senior Data Scientists (8+ YOE)	Moderate	Limited GenAI transition bandwidth at leadership levels

These roles are not easily sourced externally — and often require structured reskilling and stack-level fluency that most lateral candidates lack.

FY25 Talent Supply Gap: Quarterly Heatmap

To understand how scarcity evolved over the year, we mapped the candidate shortfall per 10 job postings across seven high-demand AI skill areas. The result is a calibrated, quarter-by-quarter heatmap showing where market pressure peaked.

Calibrated Talent Supply Gap Heatmap by Skill Cluster and Quarter (FY25)

Skill Cluster	Q1 (Apr-Jun 24)	Q2 (Jul-Sep 24)	Q3 (Oct-Dec 24)	Q4 (Jan-Mar 25)
GenAI Engineering	3.8	4.5	4	3.5
MLOps Specialists	4.2	5	4.5	4
Cloud-native Data Architects	5	5.5	5	4.5
AI Governance Specialists	6	6.5	6	5.5
Senior Data Scientists	4.5	5.2	4.7	4.3
Advanced Data Engineers	4.2	5	4.5	4
AI Product Managers	5	5.6	5	4.6

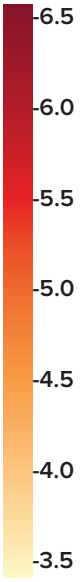

Quarter

- GenAI Engineering showed a worsening gap in Q2 (4.5), ending the year at 3.5 — still one of the most constrained roles.
- AI Governance Specialists faced the highest systemic pressure, peaking at 6.5 in Q2 and remaining elevated in Q4.
- MLOps, Cloud-native Architects, and Product Managers also faced persistent shortages, especially in the first half of FY25.
- Senior Data Scientists and Advanced Data Engineers remained in short supply throughout, reflecting the growing demand for platform maturity and domain fluency.

Scarcity and adoption dynamics also vary by industry. The following table highlights how key sectors are engaging with GenAI technologies, and the business imperatives driving this adoption:

Sector	GenAI Use Cases	Adoption Drivers
Automotive	Code generation co-pilots, ADAS testing automation	Accelerate EV and autonomous system softwaredevelopment
Healthcare	Clinical trial summarization, diagnostics, regulatory AI	Enable personalized care, reduce approval timelines
BFSI	Document automation, fraud detection, investment research	Mitigate risk, enhance customer experience, streamline compliance
Retail	Personalization engines, dynamic pricing, supply chain optimization	Scale GenAI pilots into ROI-backed systems
Energy	Predictive maintenance, real-time forecasting, ESG tools	Improve sustainability, forecasting, and operational resilience

Strategic Signals

- Hiring velocity is slowing: GenAI Engineers now take over twice as long to close as traditional ML roles.
- Mid-year pressure is real: Most roles peaked in scarcity in Q2, highlighting a recurring hiring pinch-point.
- The gap isn't closing evenly: Even as Product and Data roles ease slightly, infrastructure and governance roles remain under-filled.

Organizations cannot hire their way out of this. For roles like GenAI Engineering or LLMOps, the path forward lies in structured transitions — moving cloud, big data, and platform engineers into new GenAI-aligned roles through focused internal programs.

Summary Takeaway

FY25 was not just a competitive talent year — it revealed a structural imbalance. Scarcity is no longer isolated to niche roles. It's now embedded in the AI lifecycle itself — from build to deploy to govern. Employers must stop thinking of hiring as the only solution. Instead, they must start building talent — from within.
As the gap between talent demand and availability widens, the question is no longer if organizations should act, but how they must respond — structurally, strategically, and systemically.

Organizational Implications

If FY25 made anything clear, it's that India's GenAI talent gap is not just a sourcing issue, it's a systems issue. As the supply-demand imbalance deepens, the organizations that will win are not those with the highest hiring budgets, but those with the clearest internal answers to three questions:

- What skills do we need to own in-house, vs. rent from the market?
- Where can we cross-skill or convert existing roles into GenAI functions?
- How do we embed GenAI capabilities across our tech, product, and risk layers?

This section outlines the organizational responses that matter most in a high-scarcity, high-velocity AI adoption landscape.

Redesigning Role Architecture

The traditional boundaries between data scientists, software engineers, and cloud architects are collapsing. In their place, blended, stack-specific roles are emerging, often too new for existing job families to accommodate.

To keep up:

- Workforce planning must adopt modular role definitions, allowing skill-based transitions into emerging functions (e.g., Cloud Engineer → LLMOps Specialist)
- Org structures must absorb hybrid capability clusters (e.g., Prompt Engineering teams reporting into Platform or Product)

The new job families aren't linear — they're layered and composable. Organizations must move beyond rigid hierarchies and toward fluid, stack-informed talent models.

Internal Mobility is the New Hiring Strategy

In a high-scarcity market, buying talent externally is expensive, slow, and increasingly ineffective. Internal mobility, from adjacent roles like cloud, DevOps, data ops, or analytics — is the only scalable path forward.

Leaders must:

- Map adjacent-to-GenAI pathways for at least 20–30% of their engineering workforce
- Build role-based skilling roadmaps, tied to business-aligned capability needs (not generic certifications)
- Enable high-speed pilot transitions (e.g., hackathons, GenAI guilds, shadowing) to validate fit before full upskilling investment

Embed Learning into Engineering

Upskilling cannot sit in HR alone. In GenAI, the engineering function must own L&D accountability — because success depends on operational tool fluency, not classroom hours.

This means:

- Building stack-aligned training programs (e.g., Pinecone, LangChain, Hugging Face, LLMOps)
- Integrating training into sprint cycles (e.g., 20% time models, lab weeks, in-sprint experimentation)
- Creating role-linked skill assessments that directly impact progression and deployment eligibility

Operationalize GenAI Across the Org

As GenAI goes from prototype to platform, the capability must move beyond isolated teams. This requires:

- Cross-functional governance layers, including Legal, Risk, and Ethics
- PM-led GenAI playbooks that integrate LLM workflows into core product teams
- Shadow deployment and experimentation environments where engineers can safely test GenAI stack components

The organizational transformation is not just about upskilling engineers — it's about preparing the whole enterprise to deploy, measure, and govern GenAI at scale.

Summary Takeaway

The GenAI talent challenge is not a pipeline issue — it’s a structural challenge that requires organizations to build new talent flows, decentralize capability ownership, and treat learning as a product, not a program. Companies that treat GenAI readiness as an HR project will struggle. Those that embed it into engineering, product, and governance systems will lead.

With roles, readiness, and responses now redefined, we turn our attention to what lies ahead: the core skills and capabilities that will shape AI teams in FY26 and beyond.

FY26 Skill Priorities — What Must Be Built?

The demand for AI talent is no longer about who understands models — it's about who can make them usable, govern them responsibly, and integrate them at scale. In 2025, skill value will be defined not by breadth, but by business proximity. The most sought-after capabilities will be those that accelerate GenAI productization, enable risk-aware deployment, and support AI-driven decision-making across the enterprise.

This section consolidates the most critical skills for FY26, grouped into strategic clusters. These are based on projected demand intensity, scarcity signals, and enterprise adoption maturity — not just keyword frequency in job postings.

The Strategic Skill Stack for FY26

Skill Area	Why It Matters in FY26
Prompt Engineering	Evolves from role into a team-level competency for GenAI UX and performance tuning
LLMOps / AI Lifecycle Mgmt.	Anchors model orchestration, logging, feedback loops, and continuous retraining
RLHF & Instruction-Tuning	Becomes critical to custom LLM optimization, especially in enterprise-facing chat systems
Vector Database & Embedding Ops	Underpins retrieval-augmented GenAI (RAG) — foundational for scalable LLM integration
Advanced Data Engineering	Pipelines now demand orchestration readiness (e.g., Airflow, Dagster), vector DB fluency (e.g., Pinecone, Weaviate, Qdrant), and privacy-aware processing
Cloud-Native GenAI Infrastructure	Involves GPU scaling, Bedrock, Vertex AI, Azure OpenAI — now core to model operations
AI Product Management (GenAI)	Shifts from MVP oversight to integration of GenAI into roadmaps, workflows, and tooling
AI Governance & Risk	Cross-functional, sector-specific skillset tied to compliance, ethics, and risk mitigation
Model Evaluation & Explainability	Gains priority as LLM performance becomes business-critical; integrates QA + XAI
LangChain / OSS Stack Fluency	Hands-on familiarity with open-source libraries (e.g., LangChain, LlamaIndex, Hugging Face)

These capabilities are not just hiring priorities — they are preconditions to GenAI production-readiness.

Emerging Enablers & Sector-Specific Priorities

Specialized Skill Area	Application Context
Domain-tuned LLMs	Sector-specific LLM fine-tuning (e.g., BFSI, legal, healthcare)
AI-Augmented Analytics (NLQ/XAI)	Bridges GenAI with BI and self-service decision support
Retrieval-Augmented Generation (RAG)	Dominant GenAI deployment pattern for knowledge-intensive use cases
Simulation & Scenario Modelling	Critical in BFSI and operations-heavy sectors for AI-driven forecasting
AI-as-a-Service (AaaS) Platform Fluency	Bedrock, Vertex AI, Model Garden — enabling faster deployment at scale

These skills may not appear in every job spec, but they will increasingly define which teams ship on time, remain compliant, and unlock GenAI at scale.

What's Fading, What's Flattening

While most AI fundamentals will remain relevant, their differentiating value is declining. The following skills are being abstracted, commoditized, or absorbed into platforms:

Skill Area	FY26 Outlook
BI Reporting Tools	Still essential for monitoring, but increasingly separate from AI/ML workflows
Generic ML Algorithms	Abstracted via APIs; no longer core to hiring decisions
NLP Preprocessing	Replaced by foundation model embeddings and GenAI pipeline frameworks
Basic Cloud Certifications	Insufficient alone to assess AI deployment readiness

Strategic Skill Development Imperatives

Organizations that intend to lead on GenAI must recalibrate their talent systems around:

- Use case–backed skill development, not open-ended L&D catalogs
- Cross-functional fluency, GenAI deployment now blends engineering, product, infra, and governance
- Open-source tooling comfort, with LLM stacks evolving weekly, agility trumps platform loyalty

In FY26, the edge will belong to teams that treat skills as systems, not siloed learnings.

Summary Takeaway

FY26's most critical AI skills won't be found in generic job descriptions — they'll be embedded in platform-readiness, orchestration fluency, and ethical scaling. The employers who win will be those who identify, nurture, and deploy talent around these capabilities — and those who understand that the real war for talent is no longer about hiring... it's about activation.

As we close, we synthesize what all this means for organizational planning — and outline the strategic priorities for talent leaders, COOs, and capability architects in the AI-first era.

Outlook & Strategic Priorities

India's GenAI talent landscape is entering a defining phase. The era of experimentation is giving way to scale — and that scale is being shaped not just by tools, but by talent. The key question facing leaders in 2025 is no longer *Can we build with GenAI? — but rather, Can we do it at speed, at scale, and with the right capability mix to sustain value?*

This final section outlines the strategic priorities that will define how organizations hire, build, and mobilize AI talent in the year ahead.

From Role-Filling to Capability Architecture

AI hiring in 2025 is not about roles. It's about readiness. Capability-based planning has moved from strategic ideal to operational necessity. Organizations must:

- Move beyond hiring for generic AI roles and toward stack-informed talent models
- Anchor team design around product-critical capabilities (e.g., vector search, GPU orchestration, RAG pipelines)
- Define internal pathways to grow GenAI fluency from adjacent functions (cloud, data engineering, product)

The best talent strategies in 2025 won't be reactive. They'll be architected.

Redesign the Talent Supply Chain

Scarcity is now structural — not cyclical. GenAI infrastructure and governance roles will remain hard to hire and slow to onboard. Organizations must:

- Invest in capability-building labs to pilot skills in context (LangChain, Pinecone, RLHF tuning)
- Deploy high-intensity upskilling tracks for platform engineers, not broad-based L&D
- Create talent mobility blueprints to shift engineers from legacy roles into next-gen stacks

The future supply chain for AI talent will be built as much inside the organization as outside it.

Location and Stack Strategy Must Converge

Talent planning must now align where you hire with what you build. This means:

- Deploying core GenAI product and infra teams in Tier-1 cities with deep stack maturity (e.g., Bengaluru, Hyderabad)
- Scaling orchestration, analytics, and RAG-support roles in Tier-2 hubs through focused training and infrastructure investment
- Shaping internal delivery models around skill concentration, not geography alone

Governance is No Longer a Compliance Problem

As GenAI enters production environments, governance is becoming a business capability. Talent strategies must:

- Integrate legal, risk, and product into GenAI planning from the start
- Invest in roles like AI Governance Consultants, Risk Modelers, and Regulatory AI Specialists
- Prepare for increased scrutiny in BFSI, healthcare, and global GCCs subject to new AI acts

AI fluency in 2025 must be technical and ethical — because product risk is no longer just a tech concern.

Measure What Matters

What gets measured, gets scaled. The new metrics for GenAI capability health must include:

- Production velocity: Time from prototype to first deploy
- Readiness coverage: % of teams with stack fluency in GenAI-enabling tools
- Internal mobility rate: Conversions into GenAI roles from adjacent pipelines
- Scarcity mitigation ROI: Time and cost saved via internal reskilling vs. external hiring

Final Takeaway

India's position in the global AI economy will not be determined by the number of engineers it produces — but by how fast it can activate, specialize, and deploy them in GenAI-ready roles. For organizations, this means shifting from talent acquisition to talent orchestration — treating AI capability as a system to be designed, not a market to be chased.

The winners of 2025 will be those who understand that AI talent isn't just scarce — it's strategic. And that building it is the new competitive edge.

Disclaimer

This document constitutes marketing material and has been prepared by Quess Corp Limited ("Quess Corp") for informational and illustrative purposes only. The content herein is based on internal data and other sources believed to be reliable; however, such information has not been independently verified. While every e ort has been made to ensure accuracy and completeness, Quess Corp makes no representations or warranties, express or implied, as to the accuracy, completeness, or reliability of the information contained in this document.

The views and opinions expressed are those of Quess Corp as of the date of publication and are subject to change without notice. Quess Corp has no obligation to update or revise this material, nor does it assume any responsibility or liability for any error, omission, or inaccuracy.

This document may contain forward-looking statements, which include words such as "expects", "plans", "will", "estimates", "projects", or similar expressions. These statements are based on current expectations, estimates, and projections, and are subject to risks, uncertainties, and assumptions. Actual results may di er materially from those expressed or implied in such forward-looking statements.

Nothing in this document constitutes or should be construed as legal, financial, investment, tax, or other professional advice or a recommendation of any kind. The content is intended solely for the recipient and may not be copied, distributed, or reproduced in whole or in part without the prior written consent of Quess Corp.

QUESS
WINNING TOGETHER

Quess Corp Limited

Sky Walk Avenue, 32/4, Hosur Rd,
Roopena Agrahara, Bengaluru,
Karnataka 560068

Follow us on