DIGITALFABRIC® AISWITCH™

prosight™
AI Engineering Capabilities Study

Series: AI-Storytelling

# Building a Future-Ready Agentic & GenAI Centre of Excellence (CoE) for agentic-generative AI CoE

Focus Area: AI
Document Type: Knowledge Article

www.digitalfabric.group

# Introduction:

AI must be a board-level mandate not managed as just a technical upgrade, but as a path to strategic transformation. As CEOs commit to an AI-first vision, the next critical question emerges:

How can large enterprises operationalize this vision with scale, speed, and safety?

The answer lies in creating an agentic-generative AI Centre of Excellence (CoE) not as a side project or innovation lab, but as a cross-functional operating system embedded into the core of enterprise execution.

Based on the CoE organizational blueprint by AISWITCH.org, this blog outlines how to design a scalable and future-ready Agentic-Generative AI CoE spanning strategy, execution, governance, and external collaboration.

Here is a level-by-level look at the roles and their responsibilities in a suggested structure for an Agentic-Generative AI Centre of Excellence.

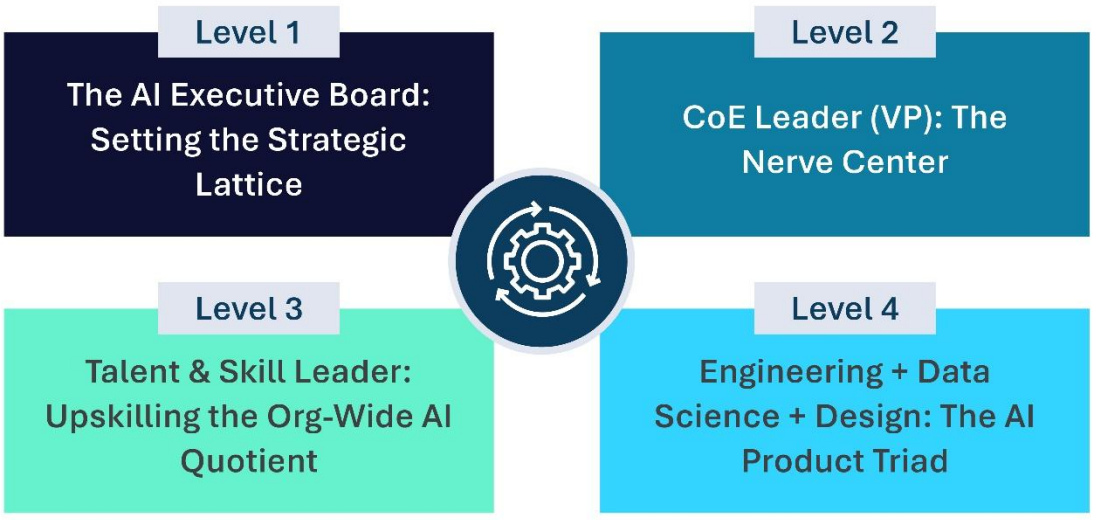


## Level 1

**The AI Executive Board: Setting the Strategic Lattice**

At the top of the structure sits the AI Executive Board of Governance, which acts as the strategic core of the CoE. This board goes beyond providing oversight, by defining the prioritization logic for use cases, aligning use cases with enterprise-wide KPIs (e.g. cost reduction, CX improvement, creating new revenue streams), and ensuring that all AI efforts are tied to business value.

The executive board isn't a standalone entity. It feeds directly into the CoE Leader (VP), who is tasked with operationalizing the goals. The handoff should not be top-down, but cyclical. The CoE Leader sends progress, risk and ROI feedback up to the board, completing the loop and enabling continuous recalibration of AI priorities.

**Key Outcomes:** Strategic alignment, prioritization logic, closed-loop governance.

## Level 2

**CoE Leader (VP): The Nerve Center**

The CoE Leader (typically a VP) is the functional integrator who ensures that GenAI isn't trapped in proof-of-concept (PoC) purgatory. Reporting into the Executive Board, this role manages:

- Implementation of GenAI projects, ranging from document summarization to customer support chatbots to GenAI-powered ERP assistants.
- Regular engagement with the enterprise's legal, finance and HR departments, not just for compliance but to forecast AI's impact on people and processes.
- Facilitation of cross-functional orchestration among engineering, data science, design and training leaders.
- This role also anchors the governance, risk and compliance (GRC) layer, which is woven throughout the CoE not siloed.

In the hat of directly reporting to the AI Executive Board, the CoE leader also serves as their eyes and ears, in terms of scanning the constantly changing technology supply environment of AI:

- Services and technology partners ecosystem development, research progress, best practices from competition and cross-industry peers, start-up's
- AI infrastructure- green data canters, test-time and inference compute infra for machine reasoning and state-space machines
- High-performance computing, quantum computing, quantum algorithms
- Multi-agent, multimodal, multi-model AI solution architecture
- Next-gen agentic-robotic-generative techniques e.g. token formers, DTQL (Digital Twin Query Language)
- Agentic-robotic interaction design and interoperability protocols e.g. A2A/ MCP current-state, machine2machine, machine2legacy-systems, machine-human

**Key Outcomes:** Prevents model sprawl, accelerates time-to-value, orchestrates cross-functional AI delivery, scans for technology change triggers, builds ecosystems with partners/ start-up's, drives internal AI application research and innovation.

**Legal & GRC Director: Trust by Design**

As agentic-generative AI models continuously learn and reason and update themselves, they produce dynamic outputs. Hence, the risk boundaries become time-elastic and thereby harder to define in terms of threats and attacks and risk profiles. That's why the Legal & GRC Director sits within the CoE framework, not outside it.

This function directly supports the CoE Leader by:

- Reviewing regulatory exposure (e.g., IP, hallucinations, data leakage).
- Setting compliance guardrails on model usage (e.g., retrieval-enhanced generation/RAG with redacted documents, content moderation filters).
- Working with the engineering team to embed explainability and traceability into model pipelines.
- Measuring potential impacts of new threats and vulnerabilities emerging from new applications and new algorithms and data infrastructure
- Collaborating with SP/SI teams in terms of auditing their risk exposures and practices, and with IT auditors/ data and AI auditors (roles that will come soon for every end-user industry, in 1-2 years' time)

Output from this role also feeds into training design so that AI literacy will be technical and ethical.

**Key Outcomes:** Reduces regulatory, reputational and ethical risks; supports responsible AI deployment.

## Level 3

**Talent & Skill Leader: Upskilling the Org-Wide AI Quotient**

Agentic-Generative AI without human adaptation = value leakage. That's where the Talent/Training Leader becomes critical. As every business increases their data, AI, intelligence leverage, every role in their ecosystem be it employees, customers, partners, suppliers, will have a digital twin of themselves in terms of an "intelligence" profile.

These agentic twins of humans need to be managed as co-workers, so this hybrid man-machine workforce will require a drastic change in stakeholder behavior and different definitions of accountability. The talent management function will have to pivot drastically.

To start with, the talent leader within the AI CoE will collaborate with the Legal Director and the general HR function, to ensure that training programs reflect governance guidelines. S/he also works with the CoE Leader to prioritize up-skilling based on active use cases. S/he has the biggest change evangelist role within the enterprise, to impact the behavioral switch in all key stakeholders.

**Examples include:**

- Prompt/ agentic development and application engineering bootcamps for different functional teams like sales, marketing, operations, finance, HR, legal etc.
- LLMOps, ModelOps, DataOps, AgentOps training for IT and AI tech teams
- Risk-aware AI literacy for legal and compliance teams.

Training is also tailored to insights from project retrospectives—shared from engineering and design teams.

**Key Outcomes:** Builds agentic and GenAI fluency, aligns learning with real use cases, drives change evangelism @scale, accelerates adoption.

**Partner Ecosystem VP: Scaling Through Collaboration**

The Partner Ecosystem Leader stitches together external threads such as:

- AI research labs for co-developing fine-tuning strategies.
- Hyperscalers for access to APIs (OpenAI, Claude, Gemini, etc.).
- Startups for specialized plugins or domain-tuned models.
- Academia for workforce development and responsible AI frameworks.

The partner ecosystem VP's work directly informs the Engineering stack (Level 4) and enables the Training Leader to roll out certified learning paths (e.g., via Google Cloud or AWS AI certification).

**Key Outcomes:** Speeds innovation, brings in domain depth, enables global-grade scale.

## Level 4

**Engineering + Data Science + Design: The AI Product Triad**

At the technical core of the CoE are three interlinked units—Engineering, Data Science and Design. Together, they transform ideas into enterprise-grade AI products.

**Engineering Director:**

- Leads solution architecture.
- Works with data engineers to build agentic solutions and RAG pipelines for example, orchestration layers and API gateways.
- Aligns infrastructure decisions (GPU clusters, ASIC, model hosting) with usage volume forecasts from business teams.
- Future-proofs and prevents vendor lock-in

**Data Science Leader:**

- Designs model evaluation metrics that align with business KPIs.
- Partners with Engineering to build fine-tuned models and RAG stacks.
- Works with the Design Leader to fine-tune model prompts based on UX feedback.

**Design Leader:**

- This role bridges the human-AI interaction loop. Efforts here are about more than UI.
- Works on use-case specific prompt templates, chatbot flows and GenAI interaction modalities.
- Feeds user friction points to Engineering for latency, reliability or UX fixes.

Together, this triad ensures that AI solutions are usable, scalable and explainable—not just technically elegant.

**Key Outcomes:** Scalable, usable and explainable GenAI products.

**Getting More Out of the Agentic GenAI CoE Structure**

One of the advantages of using this structure is that it will continue to guide organizations as they advance through agentic generative AI maturity stages. To help enterprises benchmark their progress, here's a simplified maturity model:

| Stage | Description |
|---|---|
| 1. Awareness | PoCs and experimentation |
| 2. CoE Formation | Governance and cross-functional alignment |
| 3. Enablement | Role-based training and use-case onboarding |
| 4. Productization | In-house Agentic Generative AI IP and frameworks |
| 5. Business Model Shift | Agentic Generative AI as a core differentiator |

**Feedback Loops: The Hidden Superpower**

Each node of this CoE structure should be in constant feedback with others. Legal insights influence training. Design feedback improves engineering pipelines. Executive goals inform model evaluation metrics. This recursive structure creates a flywheel effect for GenAI growth.

**Change Management Considerations:**

Establishing an effective agentic generative AI CoE also requires strong change management. The CoE will regularly require:

- Re-skilling at scale
- Redefining performance metrics
- Breaking silos held by data, tech and business units

This is where leadership sponsorship and transparent communication make or break success.

## Conclusion: View Agentic Generative AI as an Operating System, Not a Side Project

When the Centre of Excellence is structured as described above, the organization moves beyond a taskforce approach. Agentic and generative AI is no longer treated as a pilot or an isolated initiative it becomes an operating system that spans vision, talent, data, design, compliance, and external engagement.

Every function feeds into the others, creating a closed feedback loop that enables:

- Faster time to market
- Reduced model risk
- Higher adoption across business units

In short, it transforms AI into a core competency not a one-off experiment.

# Reach out to us

**Tap into expert insights and transformative strategies to lead your enterprise into the age of AI, Cloud, and Intelligent Automation.** Connect with us to accelerate growth, optimize operations, and unlock new value at scale.

Email ID: rishant.singh@digitalfabric.in

Learn more about DIGITALFABRIC® for more insights:
https://digitalfabric.group

Connect with DIGITALFABRIC®

prosight™
AI Engineering Capabilities Study

DIGITALFABRIC® AISWITCH™