

MAY 2026

FIRST EDITION

SKILLPULSE

Exclusive Magazine for leaders building Skills First Organizations



COVER STORY Skill Cataloging is Back!

Learn how new age L&D Leaders are leveraging AI to reimagine the most boring yet most critical task in L&D – Cataloging skills that are evolving at AI pace

Q [ENQUIRIOUS.COM](https://enqurious.com)

INSIDE THIS ISSUE
SKILL ONTOLOGY ·
AI-NATIVE COHORTS
· COMMUNITY OF
LEADERS · L&D DIARY



“Skill is the new oil that deliver business outcomes via engines called Roles.

TRACE is the blueprint that helps design the engine!”

TABLE OF CONTENTS

01

FROM THE CEO DESK

Trading the Spreadsheet for a GPS: Why Your Organisation Needs a Skill Ontology Now

Amit Choudhary · Co-Founder & CEO, Enqurious

From static taxonomies to dynamic, AI-driven skill graphs. How ontologies reveal "Skill Adjacency" and enable precision learning at scale.

02

THE ARJUN COLUMN

We Don't Have a Skills Gap. We Have a Skills Visibility Problem.

Arjun Sivasundar · Chief Mentor, Enqurious

People without jobs and jobs without people. The real issue isn't a shortage of skills but an inability to see the ones that already exist.

03

CO-FOUNDER'S NOTE

Why We Built a Community of Data & AI Leaders

Prateek Kumar · Co-Founder & COO, Enqurious

Skills alone aren't enough. The combination of skills, tools, and approach, shaped through peer-level exposure, is what separates high-impact professionals.

04

FROM OUR MENTORS

Databricks Is Investing \$250 Million in India. Is Your Team Ready?

Mansi Mutreja · AI & ML Engineer, Enqurious

Roughly 20% of data professionals can navigate Databricks without help. The gap between platform capability and team readiness is where ROI goes to die.

05

FROM OUR MENTORS

How a Tiny Change Can Greatly Impact Performance: Delta Lake's Deletion Vectors Explained

Sayli Nikumbh · Contributor, Enqurious

Updating one row used to mean rewriting an entire file. Deletion vectors flip that, cutting storage footprint by up to 97% and total cost by up to 60% on update-heavy workloads.

CONTENTS CONTINUED

06

FROM OUR MENTORS

Your Data Model Is Your AI Product Now

Mandar Sawant · Contributor, Enqurious

Same dataset, same AI model, different data structure, opposite answers. Why the semantic layer, not the algorithm, determines enterprise AI success.

07

GUEST COLUMN

The End of the Static Spreadsheet: How AI Is Rewriting the Skill Catalogue

Siddharth Ghose · Brillio

Agentic AI can scrape job descriptions and project outcomes to update taxonomies in real time, freeing L&D leaders to coach rather than catalogue.

08

GUEST COLUMN

The New Career Ladder: From Linear to Modular

Manisha B. · L&D Leader, Brandon Hall Gold Award Winner

The junior-to-senior ladder is breaking. Growth now moves through modes of working, from Learner to Strategist, measured by complexity handled rather than reports managed.

09

GUEST COLUMN

Beyond the Matrix: The Rise of Generative Skill Ontologies

Tarang Vadalikar · Tiger Analytics

The skills matrix has become dangerous, not just tedious. Why the rigid taxonomy tree is giving way to a living knowledge graph that maps capability by relationships, not job titles.

FROM CEO DESK



AMIT CHOUDHARY

Founder & CEO

Trading the Spreadsheet for a GPS: Why Your Organization Needs a Skill Ontology Now

The role of a leader in today's workforce is no longer just about managing people; it is about managing the "skill velocity" of the entire organization. We are witnessing a fundamental shift from rigid job descriptions to fluid skill ontologies a transformation that is not just a trend, but a survival mechanism in an AI-driven landscape.

From Static Libraries to Dynamic GPS

For decades, organizations relied on skill taxonomies flat, manual spreadsheets that categorized people into buckets. But a taxonomy is like a static library; it tells you what you have, but not how to get where you're going.

In contrast, a skill ontology is like a GPS. It is a structured, dynamic data model that understands the complex relationships between skills, roles, and business outcomes. It recognizes that "Python" isn't just a category it's a prerequisite for "Machine Learning" and a neighbor to "DevOps". This semantic layer allows us to identify "Skill Adjacency," revealing that an analyst might only be a "20% leap" away from becoming a data scientist.

The AI Imperative

Why the urgency now? The half-life of skills is shrinking. As AI automates routine tasks, the nodes in our organizational skill graphs are shifting in real-time.

A practical, modern ontology uses **Natural Language Processing (NLP)** to scan internal data from GitHub commits to project reports to update itself automatically. This "Living" aspect ensures that the organization's understanding of its talent is never more than a few days old. We are moving toward a **Phase 3** evolution: AI-driven, predictive gap analysis that tells us what we will need in 2027 based on the trajectory of technology today.

Precision Learning and Outcome Measurement

This isn't just about documentation; it's about the **"Precision Learning"** experiences we design. By mapping learning content to specific nodes in the ontology, we can offer the "shortest path" to competency. If the system knows an employee has mastered "Excel Modeling," it won't waste their time; it will serve the exact micro-credential they need for "Strategic Financial Planning". Furthermore, we can finally measure what matters: **Skill Velocity**. Instead of tracking "course completion" (a vanity metric), we track the rate of movement from "Novice" to "Expert" across the graph. We can finally link a 10% increase in "Cloud Architecture" proficiency directly to a reduction in system downtime.

The Path Forward

As we scale ventures like Enqurious Academy, our mission is to move organizations away from "one-size-fits-all" training toward these precision ecosystems. The goal is a "Skills-First" organization where we document what a person *can do*, not just what they *are*.

In an era of rapid automation, the most valuable asset isn't a fixed set of skills; it's the ability to navigate the ever-changing map of human capability. It's time to trade your spreadsheets for a GPS.



Early in my corporate career, I worked at a services company where two problems dominated every resource planning meeting. We had people on the bench skilled professionals between projects, costing the business money. And we had unfilled project roles critical positions we couldn't fill, costing us delivery capacity and client confidence. It seemed we had to simultaneously solve two problems: **people without jobs, and jobs without people.**

I soon began to realise what was happening. The issue wasn't a shortage of skills. It was that we couldn't see the skills we had. A developer might sit idle while we searched for someone with API integration experience despite having built three integrations in his previous project. A business analyst on the bench may have facilitated stakeholder workshops for years, but we were hiring externally for a change management role because "that's a different role."

I've since seen this pattern everywhere. Most organisations believe they're short of critical skills. It's a familiar refrain in boardrooms and strategy reviews: we need more data capability, stronger commercial acumen, better digital fluency. The response is equally predictable launch a hiring drive, commission a training program, bring in consultants. Yet despite sustained investment, the gaps persist. Leaders still report that they can't staff key projects. Delivery still depends on a handful of over-stretched individuals.

The conclusion seems unavoidable: we simply don't have enough of the right people. But there's an uncomfortable alternative explanation. What if many of those skills already exist but remain invisible, untrusted, or undeployable?

We don't have a skills gap.

We have a skills visibility problem.



**ARJUN
SIVASUNDAR**

Cheif Mentor, Enqurious

Consider a mid-sized professional services firm struggling to resource a data transformation project. After weeks of searching, they hire an external analytics lead at significant cost. Three months later, a manager in operations mentions in passing that she'd built the prototype for a similar system two years earlier, and it had worked well. No one

outside her immediate team knew about it. She'd never been asked. This isn't an edge case. It's systemic. Skills exist, but decisions are made as if they don't. Leaders make talent decisions using proxies: job titles, role definitions, past performance reviews. These are useful for governance and reporting. They are poor indicators of current capability. Because skills change rapidly. A finance analyst teaches herself Python to automate a recurring report. A project manager picks up stakeholder facilitation by running a contentious procurement process. A customer service lead develops pattern recognition in complaint data that matches human accuracy. These capabilities are real. But they're rarely recorded, validated, or visible beyond their immediate context. As a result, under-utilised capability goes unnoticed.



The same people get tapped repeatedly because they're known quantities. Perceived gaps widen even as real capability grows quietly in pockets across the organisation. Most organisations are convinced they're short of critical skills. The truth is uncomfortable: in many cases, the capability already exists it's simply invisible to the people making decisions.

THREE SIGNS YOU HAVE A VISIBILITY PROBLEM, NOT A SKILLS PROBLEM

1. You consistently hire externally for roles where someone internal has done adjacent work but you didn't know.
2. The same handful of people get tapped for new initiatives, while equally capable colleagues remain on the bench.
3. Capability conversations happen by job title and role definition, not by what people have actually built or done.

“What looks like a skills shortage is often a visibility failure.”

The capability exists. The signal doesn't.



PRATEEK KUMAR

Co-Founder and COO

Why We Built a Community of Data & AI Leaders

When we first stepped into the world of data and AI upskilling, the problem didn't appear to be a lack of content. There were courses, certifications, bootcamps, and endless learning pathways. On the surface, the ecosystem looked rich. But as we worked closely with enterprises and professionals, a deeper gap became evident the absence of meaningful connection between learning, real-world application, and leadership thinking. That realization is what led us to initiate a community of Data & AI leaders.

The modern data professional is not just expected to write code or build pipelines. They are expected to influence business decisions, design scalable architectures, and continuously adapt to an evolving technology landscape.

Yet, most learning environments operate in silos individuals learning alone, teams solving problems in isolation, and organizations reinventing the wheel repeatedly. We believed there had to be a better way. This community was not built as another networking group or a broadcast platform for content. It was envisioned as a living ecosystem a space where leaders could exchange ideas, challenge perspectives, and learn from each other's real experiences. Because in the world of data and AI, context matters as much as capability. decisions, design scalable architectures, and continuously adapt to an evolving technology landscape. They are expected to influence business decisions, design scalable architectures, and continuously adapt to an evolving technology landscape. Yet, most learning environments operate in silos individuals learning alone, teams solving problems in isolation, and organizations reinventing the wheel repeatedly. We believed there had to be a better way. This community was not built as another networking group or a broadcast platform for content.



It was envisioned as a living ecosystem a space where leaders could exchange ideas, challenge perspectives, and learn from each other's real experiences. Because in the world of data and AI, context matters as much as capability. decisions, design scalable architectures, and continuously adapt to an evolving technology landscape. Yet, most learning environments operate in silos individuals learning alone, teams solving problems in isolation, and organizations reinventing the wheel repeatedly. We believed there had to be a better way. This community was not built as another networking group or a broadcast platform for content. It was envisioned as a living ecosystem a space where leaders could exchange ideas, challenge perspectives, and learn from each other's real experiences. Because in the world of data and AI, context matters as much as capability.

One of the biggest insights we gained early on was this: skills alone are not enough. What truly differentiates high-impact professionals is the combination of skills, tools, and approach. While skills can be taught and tools can be learned, approach is shaped through exposure exposure to real problems, diverse thinking, and peer-level conversations. That is exactly what a strong community enables.

Because in the end, the future of Data & AI will not be shaped by individuals working in isolation, but by communities growing together.

From a founder's standpoint, this initiative is also about shifting the narrative of upskilling. Traditionally, learning has been treated as a transactional activity enroll, complete, certify. But the future of learning is continuous, collaborative, and deeply integrated with work. Communities play a critical role in enabling this shift by making learning more dynamic and context-driven.

Another important reason behind building this community is to bridge the gap between industry and talent. Organizations are constantly looking for professionals who can deliver impact from day one, while professionals are seeking clarity on what "industry-ready" truly means. By bringing leaders, practitioners, and learners onto the same platform, we create a shared understanding of expectations, challenges, and opportunities.

We are also seeing the rise of AI transforming not just technology, but the way we learn and work. In such a fast-changing environment, no single organization or individual has all the answers. Progress depends on collective intelligence on how effectively we learn from each other.

This community is our step toward enabling that collective intelligence.

It is a space for honest conversations, for practical insights, and for building a culture of continuous evolution. As founders, we see this not just as an initiative, but as a responsibility to contribute to an ecosystem where learning is not just consumed, but co-created.

HIGHLIGHTS

Nearly 80% of a data professional's time is spent not on AI models, but on cleaning and preparing data. The real superpower isn't just building algorithms it's asking the right questions and making messy data actually usable.

FROM OUR MENTORS

"Databricks can now do everything. I just came from a session where they were showing how to build an app from scratch using Genie Code it took 45 minutes. The entire ETL pipeline, the Lakehouse, the app. All of it. I'm genuinely worried: will AI take our jobs? Do we need to go back to farming?" Partner Solutions Architect, Databricks India

The problem here isn't one that lives inside a Databricks notebook. It lives inside the gap between what the platform can do and what most professionals actually know how to do with it.

Let's start with the headline number: **Databricks is investing over \$250 million (approximately ₹2,134 crore) in India** over the next three years, beginning 2025–26. The investment is designed to expand AI and data operations, fund local R&D, launch an AI academy, grow its India headcount by over 50% (crossing 750 employees), and train **500,000 professionals** in data and AI skills.

Half a million people. That's not a learning initiative. That's an infrastructure project.

Now zoom out for a second. Walk into any data team in India in a bank, a consulting firm, a startup and ask ten Data Engineers and ten AI/ML engineers to build an end-to-end pipeline on Databricks. Ask them to set up a workspace, manage permissions, configure Unity Catalog, and get a Lakehouse running. Based on interactions with 250+ data professionals across the country, a generous estimate is that **roughly 20% would feel truly comfortable** navigating the platform without external help.

MANSI MUTREJA



What Is a Capability Gap, Exactly?

It's not about intelligence or effort. It's about misalignment between what a tool demands and what the person operating it can deliver.

Think of it this way: if a world-class architect designs a bridge but the construction crew has only ever built wooden sheds, the bridge doesn't get built. Not because the design was bad. Not because the crew was incompetent. But because there's a gap between the knowledge the project requires and the knowledge in the room.

In the data world, this plays out constantly. Organizations license Databricks, run the kick-off call, and then quietly discover that the engineers on their teams talented, experienced people are more comfortable with tools they've used for years than with a platform they've only just met. The licenses sit underutilized. The ROI conversation gets awkward. And somewhere in a finance spreadsheet, a very large number starts looking suspicious.

FROM OUR MENTORS

This isn't a criticism. It's a pattern. And recognizing it is the first step toward fixing it. "Databricks can now do everything. I just came from a session where they were showing how to build an app from scratch using Genie Code it took 45 minutes. The entire ETL pipeline, the Lakehouse, the app. All of it. I'm genuinely worried: will AI take our jobs? Do we need to go back to farming?" Partner Solutions Architect, Databricks India

The problem here isn't one that lives inside a Databricks notebook. It lives inside the gap between what the platform can do and what most professionals actually know how to do with it.

Let's start with the headline number: **Databricks is investing over \$250 million (approximately ₹2,134 crore) in India** over the next three years, beginning 2025–26. The investment is designed to expand AI and data operations, fund local R&D, launch an AI academy, grow its India headcount by over 50% (crossing 750 employees), and train **500,000 professionals** in data and AI skills.

Half a million people. That's not a learning initiative. That's an infrastructure project. Now zoom out for a second. Walk into any data team in India in a bank, a consulting firm, a startup and ask ten Data Engineers and ten AI/ML engineers to build an end-to-end pipeline on Databricks.

What Is a Capability Gap, Exactly?

It's not about intelligence or effort. It's about misalignment between what a tool demands and what the person operating it can deliver. Think of it this way: if a world-class architect designs a bridge but the construction crew has only ever built wooden sheds, the bridge doesn't get built. Not because the design was bad. Not because the crew was incompetent. But because there's a gap between the knowledge the project requires and the knowledge in the room. In the data world, this plays out constantly.



Organizations license Databricks, run the kick-off call, and then quietly discover that the engineers on their teams talented, experienced people are more comfortable with tools they've used for years than with a platform they've only just met. The licenses sit underutilized. The ROI conversation gets awkward. And somewhere in a finance spreadsheet, a very large number starts looking suspicious. This isn't a criticism. It's a pattern. And recognizing it is the first step toward fixing it.

FROM OUR MENTORS

How a Tiny Change Can Greatly Impact Performance – Delta Lake's Deletion Vectors Explained

Picture this: You have a 1GB Parquet file in Delta Lake. You need to update just one value in one row. What happens next?

Before deletion vectors, Delta would rewrite the entire 1GB file just to change one cell. It's like copying and rewriting a whole Excel sheet just to edit a single cell. Inefficient, right?

That's exactly the problem Delta Lake solved with deletion vectors, and understanding how they work is fundamental to optimizing your data pipeline performance.

WHAT ARE DELETION VECTORS?

Deletion vectors are small binary files (.bin) that logically mark rows as deleted without rewriting the original Parquet files. Instead of recreating large data files, Delta Lake creates lightweight metadata files that track which rows have been removed.

Think of it like this: instead of erasing data and reprinting the page, you add a bookmark that says "ignore rows 42–67 when reading." The original data stays intact; you just track what's no longer valid.

This becomes critical when working with large tables where even tiny updates add up. A 10GB customer table updated daily? Without deletion vectors, you'd rewrite 10GB of data every single day. With them? You write a few kilobytes of metadata.



**SAYLI
NIKUMBH**

THE MODERN APPROACH: WITH DELETION VECTORS

With deletion vectors (enabled by default in Databricks Delta Lake 3.0+), the workflow is fundamentally different. How updates and deletes work:

1. When you delete or update rows, Delta leaves the original Parquet file untouched
 2. Delta creates a small deletion vector file (.bin)
 3. This file stores bitmap information about which row IDs are logically deleted
 4. The original Parquet file remains intact
 5. During reads, Delta merges the Parquet data with deletion vectors, filtering out deleted rows on-the-fly
- The efficiency gains:
- Minimal writes: Only tiny metadata files are created (kilobytes instead of gigabytes)
 - Original files preserved: Parquet files stay intact, enabling efficient time travel and versioning
 - Faster operations: Updates complete in milliseconds rather than minutes
 - Storage efficiency: No unnecessary file duplication across versions

Practical example: That same 500M-row customer table update of 10K rows now writes ~50KB of deletion vector metadata instead of 100GB+ of new Parquet files.

FROM OUR MENTORS

REAL-WORLD IMPACT: STORAGE, PERFORMANCE, AND COSTS

The benefits become significant at scale.

Storage Impact

- Without deletion vectors: Each update generates a new Parquet file. A 10GB table updated daily becomes 300GB+ in a month.
- With deletion vectors: Same scenario uses approximately 10GB of storage + ~100MB of metadata.
- Net result: 97% reduction in storage footprint.

Read Performance


- Queries no longer search through multiple file versions for the current state
- Delta's query optimizer can skip unchanged files entirely
- Scans require fewer I/O operations
- Performance gain: 10-40% faster queries on frequently-updated tables.

Cost Reduction

- Less data rewritten means fewer compute seconds billed
- Smaller storage footprint reduces cloud storage costs
- Faster operations translate to shorter job runtimes
- Overall impact: 30-60% reduction in total cost of ownership for update-heavy workloads.

Half a million people. That's not a learning initiative. That's an infrastructure project. Now zoom out for a second. Walk into any data team in India in a bank, a consulting firm, a startup and ask ten Data Engineers and ten AI/ML engineers to build an end-to-end pipeline on Databricks.

It's not about intelligence or effort. It's about misalignment between what a tool demands and what the person operating it can deliver. Think of it this way: if a world-class architect designs a bridge but the construction crew has only ever built wooden sheds, the bridge doesn't get built. Not because the design was bad. Not because the crew was incompetent. But because there's a gap between the knowledge the project requires and the knowledge in the room. In the data world, this plays out constantly.



Organizations license Databricks, run the kick-off call, and then quietly discover that the engineers on their teams talented, experienced people are more comfortable with tools they've used for years than with a platform they've only just met. The licenses sit underutilized. The ROI conversation gets awkward. And somewhere in a finance spreadsheet, a very large number starts looking suspicious. This isn't a criticism. It's a pattern. And recognizing it is the first step toward fixing it.

FROM OUR MENTORS

Your Data Model Is Your AI Product Now

Most companies think their AI problem starts with choosing the right AI model.

Should we use OpenAI, Anthropic, or Snowflake Cortex? Should we build or buy? How quickly can we show ROI?

These are the questions dominating boardrooms.

They are also usually the wrong questions. Because the biggest reason enterprise AI fails is rarely the AI model.

It is the data model underneath it.

Many organisations are trying to build intelligent AI systems on top of data foundations they would not trust for a basic dashboard. They expect reliable insights from inconsistent dimensions, duplicated logic, and poorly structured tables. AI does not fix that. It exposes it.

And worse, it does so with confidence. That is the real risk.

The Illusion of “AI-Ready” Data

Snowflake’s Cortex platform is a strong example of where the industry is heading.

The promise is compelling: bring AI directly to the warehouse, let users ask natural language questions, and generate insights without moving data across systems.

But Cortex is only as smart as the data structure it is forced to reason against.

Give it a clean star schema with clear relationships and consistent dimensions, and the AI model behaves like a strong analyst.

MANDAR SAWANT



Give it a flat, denormalised table full of ambiguity, and it will still answer confidently. It will just be wrong. This is where many leadership teams misunderstand AI.

They assume the AI model creates intelligence. In reality, the AI model amplifies whatever intelligence already exists in the data model. If your warehouse is confused, your AI will only spread that confusion persuasively.

The Experiment

To test this, I used a supply chain dataset with 10,000 order records containing supplier details, product categories, shipment modes, order values, and delivery delays.

The same dataset was exposed to Snowflake Cortex Analyst in two ways.

First, as a single flat table where everything was collapsed into one wide structure.

Second, as a proper star schema with a central fact table and separate dimensions for suppliers, products, and dates. Then I asked both versions the same business questions.

FROM OUR MENTORS

Question 1: Which supplier region had the highest average delivery delay?

The flat data model said Europe. The star schema said North America. Same dataset. Opposite answer. The reason was simple: supplier region was inconsistent in the flat table. The same supplier appeared under multiple regions across different rows. In the star schema, supplier region came from a single authoritative supplier dimension where each supplier had one correct classification. One invisible join created the difference between a wrong answer and a reliable one. That was not an AI problem. It was a data modelling problem.

Question 2 and 3: Revenue and delayed shipment analysis

The same pattern appeared when analysing Electronics revenue and delayed shipments by shipping mode. The flat data model significantly understated results compared to the star schema, enough to materially change procurement, investment, and operational decisions. The issue again was structural. When product categories are inconsistently assigned in a flat table, the AI model reasons accurately against broken logic. It does not hallucinate. It faithfully scales the error. That should worry every enterprise leader calling their platform "AI-ready."

Your Semantic Layer Is Product Design

One of the most useful parts of the experiment came while building the

semantic layer.

Snowflake's Cortex Autopilot automatically generated business-readable descriptions for the star schema tables. It correctly understood supplier and product dimensions because the structure was clear: narrow scope, clean naming, and explicit purpose. The flat data model offered none of that. Just columns and ambiguity. This matters because a good data model does not just improve performance. It teaches the AI model how your business works. Your joins are instructions. Your dimensions are definitions. Your semantic layer is product design. If the structure is unclear, the AI does not become creative. It becomes dangerous.

The Real Risk of Enterprise AI

For years, data engineering has focused on speed: faster pipelines, larger datasets, more dashboards. Modelling discipline was often treated as something to improve later. Later has arrived. When a human analyst queries bad data, strange numbers trigger suspicion. When an AI model queries the same data, it produces a polished paragraph with confidence and perfect grammar. The error disappears behind presentation quality. That is the real danger.

Not replacement. False authority. AI gives a professional voice to technical debt. Every shortcut in modelling, every undocumented relationship, every careless dimension becomes a business risk because it sounds trustworthy while being wrong.

**The AI model was never the real product.
The data model is.**



SIDDHARTH GHOSE

Brillio

The End of the Static Spreadsheet: How AI is Rewriting the Skill Catalogue

If you work in talent development, you've probably felt the frustration. The skill catalogue is supposed to be your organisation's go-to map of capabilities, but in practice, it's just a static spreadsheet or a rigid HRIS module that's outdated almost immediately. You can spend months perfecting the taxonomy, only to see it become obsolete as the market changes.

The Anatomy of the Current Approach

Right now, the gold standard for cataloguing skills relies on a highly structured, logical architecture. We build our taxonomies from the ground up: starting broad at the Domain level, narrowing down to the Capability Family, segmenting into a Skill Group, and finally identifying the specific Skill. It's a beautiful,

necessary hierarchy designed to support real-time dashboards, learning paths, and performance measurement. Yet, large organisations rarely have just one learning path and balance multiple, distinct schools of learning. Attempting to manually map these different domains leads to inevitable silos. It's like painting a massive bridge by the time you finish with one area, another has already changed, requiring you to start over.

The Breaking Point

The fundamental challenge we face today is scale

mixed with velocity. This is especially true when your focus leans toward technical L&D rather than behavioural training. The decay rate of technical capabilities is exponential.

There are tens of thousands of specialised skills in the market right now. A few years ago, "data analysis" might have been a sufficient tag for an employee. Today, you need to know if they are proficient in Python, SQL, specific data visualisation tools, or predictive modelling. The sheer volume of data required to maintain an enterprise-level skill kit is crushing.

Furthermore, we are relying almost entirely on self-reporting or manager appraisals to update these catalogues. It becomes an administrative burden that pulls focus away from talent development and into endless data entry. We are trying to track a moving target with a Polaroid camera, capturing a static snapshot of an employee's capability that is instantly obsolete. Enter the AI-Native Catalogue. This is where AI moves from hype to something genuinely useful. We aren't just talking about chatbots suggesting generic courses; we are looking at an entirely AI-native approach to how we structure and track talent data.

The potential impact of generative and agentic AI on skill cataloguing is profound because it solves the velocity problem. Instead of a human manually mapping a new programming language into the correct capability family, agentic AI can autonomously scrape job descriptions, project outcomes, and market trends to update the taxonomy in real time.

Imagine a centralised dashboard that unifies all professional certifications, learning metrics, and academy performance indicators into a single platform. Instead of manual inputs, AI sits in the background, looking at an employee's actual work output to infer their proficiency. It understands complex skill adjacencies natively. If a platform knows an employee is an expert in one specific cloud architecture, it can automatically map the adjacent capabilities they likely possess or highlight the exact micro-skills they need to bridge a gap.

Reclaiming the Role of the Coach

Let's talk about the real-world impact on the people running these teams. Think about your last performance appraisal cycle. How much time was spent just trying to accurately document "achievements above and beyond" to justify a promotion during an HR calibration meeting?




Or consider the reality of onboarding a new team member when business is moving so fast that a formal handover is practically impossible.

In those critical moments, talent leaders and managers do not need another blank competency matrix to fill out; they need time to provide structured, continuous guidance and hands-on coaching.

Instead of chasing down employees to update their profiles, we can spend that time doing what moves the needle: establishing focused learning partnerships across different regional teams, building targeted training for interview panellists, or having meaningful career conversations. AI takes over the cataloguing, so humans can return to coaching.

From Static Ledger to Living Engine

The future of skill cataloguing isn't about building a better spreadsheet; it's about abandoning the spreadsheet altogether. By leveraging AI, we can transition our taxonomies from static, historical ledgers into dynamic, living engines. Ultimately, the deeper we go into this dynamic data, the more we realise we're not just analysing numbers and technical proficiencies we're analysing ourselves and the living DNA of our organisations. For L&D and talent teams, embracing this AI-native approach means finally getting out of the data-entry business and stepping into our true roles: building the targeted, impactful learning experiences that drive the business forward.



MANISHA B.

L&D Leader, Brandon Hall Gold Award Winner | GenAI for HR & L&D | Strategic TNI & MDP

From Linear Progression to Capability Expansion

The traditional ladder assumes a simple equation: more experience leads to promotion, and promotion leads to leadership. In practice, AI breaks this link. A professional using AI well can outperform a more senior role constrained by outdated ways of working. Similarly, leadership impact is increasingly possible without large teams. The AI era career path reflects this shift. Instead of moving "up," individuals move through stages of capability, each defined by how they engage with AI and complexity.

The New Career Ladder: From Linear to Modular

In the age of AI, careers are no longer defined by steady upward movement through fixed roles. The traditional ladder, junior to senior to manager to executive, was built for an era where work was stable, expertise deepened slowly, and authority came from hierarchy. AI disrupts all three assumptions.

Careers today are becoming modular and skill based. Growth happens by expanding what a person can handle, more complex problems, broader systems, higher stakes decisions, not simply by gaining authority over others.

The AI Era Career Path

The AI era career ladder moves from Learner to Practitioner, then to Integrator, Innovator, and finally Strategist or Leader. These are not ranks or titles. They are modes of working.

A Learner builds fluency and understanding, developing the ability to work with AI responsibly and effectively. A Practitioner begins to translate that understanding into consistent performance gains, using AI to augment everyday work.

An Integrator goes further by embedding AI into real workflows and systems, combining domain expertise with new tools to create scale. Innovators use

AI to reimagine how work is done altogether, designing new products, services, or operating models. Strategists and leaders operate at the system level, deciding not just how to use AI, but where, why, and when not to.

Importantly, a person may occupy different stages simultaneously across different skills. Someone may be a Strategist in their functional domain and a Learner in an adjacent one. This makes career growth non linear but intentional.

Why This Is a Modular Ladder

This model is modular because careers are built by stacking skills, not replacing roles. People grow by increasing their cognitive reach, judgment, and ability to integrate technology into human systems.

Movement is no longer only vertical. It can be lateral, diagonal, or cyclical. What matters is the complexity and impact of problems handled, not the number of direct reports or years served.

McKinsey's research underscores the scale of this shift. Up to 30% of all work activities could be automated by 2030, while more than 60% of roles will change significantly in skill composition rather than disappear. At the same time, demand for technological skills is projected to rise by 55%, and demand for social, emotional, and higher order cognitive skills by over 20%. Value is shifting from execution to judgment and integration.

What This Means for Organisations



For organisations, this shift requires a fundamental rethinking of how careers are designed and rewarded.

Static ladders tied to job titles must give way to skill based frameworks that make capability visible and growth tangible. Progression should reflect increasing mastery, scope, and impact rather than position alone.

Learning and experimentation must become legitimate forms of performance. In an AI driven environment, the ability to learn quickly and adapt responsibly is as valuable as execution. Organisations that reward only delivery will fall behind those that reward evolution.

Leadership also needs to be redefined. Managing people is no longer the only, or even primary, indicator of leadership. Judgment, systems thinking, ethical responsibility, and the ability to guide human and AI collaboration are becoming central leadership qualities.

Finally, AI fluency must exist at every level. AI cannot remain the domain of specialists or centralised teams. When leaders, managers, and individual contributors share a common understanding of AI's capabilities and limits, organisations make better decisions and avoid both over reliance and fear driven resistance.

Problem with "Skill Matrix"

If you walk into any HR department today, you'll likely find a ghost in the machine: the Skills Matrix. It's a comfortable, familiar spreadsheet. A neat grid of names and checkboxes designed to convince leadership that they have a handle on what their people can actually do. But the reality is that those spreadsheets have become more than just tedious; they're dangerous. We've hit a point where the shelf-life of a hard technical skill is decaying faster than the time it takes to hire for it.

Beyond the Matrix: The Rise of Generative Skill Ontologies

The traditional approach to Taxonomy was a tree. It was rigid and hierarchical. You were a "Software Engineer," which branched into "Java," which branched into "Spring Boot." It was orderly, but it was also fragile. Measurement was constrained to some form of assessments, quizzes or exams, with it having limited bearing on the actual assimilation.

Generative AI has changed the fundamental "math" of productivity. We're seeing a 1:5 efficiency factor in AI-assisted roles. One person isn't just "doing" Java anymore; they are orchestrating a suite of agents to build, test, and deploy at a scale that makes old job descriptions look quaint. When a job role mutates every six months, a static list becomes a mere shadow. It cannot capture the "liminal" skills the stuff people learn on the fly to bridge the gap between what they knew yesterday and what the AI is doing today.



TARANG VADALIKAR

Tiger Analytics

From Trees to Webs

This is why the shift toward Skill Ontology, which is more of a neural network of what the person's capabilities and readiness are, is the logical next step. More importantly, it doesn't care about rigid frameworks, or static data ; it cares about their relationships.

An AI-powered skill ontology doesn't just see that you know Python. It uses Natural Language Processing to look at the "connective tissue" of your work. For a young professional, this ontology acts as a career GPS. It realizes that "Coding" is becoming "Editing," and "Writing" is becoming "Curation." It suggests pivots in real-time, moving people across the organization not based on their job title, but on the invisible sparks of capability the AI detects in their daily output. Additionally, it can weave in the domain & functional skills as well. These are not embedded in the classical approach of segregation based on level, grade or role, but a real-time, day to day deliverables of the individual. Thus, the Gen AI agent overseeing the skills is the parrot on the captain's shoulder! The parrot can be a surveillance mechanism, a knowledge repository, a brainstorming buddy, and in certain situations even a coach! The Web of skill ontology is the knowledge graph of this Parrot!

New "Junior" Reality

Historically, a junior spent their first two years in the trenches of syntax and basic documentation. Not anymore. Today, AI handles the "grunt work" of execution perfectly. This has created a massive challenge for young professionals: they are now required to skill up to an Intermediate level on Day One.

To be an FDE, you must possess a hybrid of technical fluency, deep domain knowledge, and functional strategy, a trifecta that traditionally took years to bake. This is where the Generative AI based web of Skill Ontology becomes the young worker's greatest asset. Instead of wondering how to gain "3 years of experience" in six months, the ontology maps out the Skill Adjacencies.

It shows a youngster that if they have mastered "Logic & Prompting" (Technical) and "Stakeholder Communication" (Functional), they are only a few nodes away from "Systems Integration" (Domain)

. It turns the nebulous "intermediate level" into a visible, reachable set of connected nodes. The "youngsters" today can be "conductors in training" from the start not just writing code, but deploying it into complex business environments.

The ontology identifies the invisible bridges between what a worker does now and what they could do tomorrow. Most importantly, it can customize the journey of each learner. Take a fresh graduate with a background in data science. In a traditional taxonomy, they are stuck in a "Junior Data Scientist" bucket. In a 2026 Skill Ontology, the system recognizes that their structural thinking is 85% adjacent to "Agentic Workflow Design." By highlighting these overlaps, the ontology allows young workers to "hack" the experience gap, moving into functional roles that previously required years of tenure. This will be different for a graduate with a computer science major or someone with a MBA! The common goal, groom them to be strategists who understand the business impact of the code the AI just spit out.

But it's still a TOOL!

Conclusion: The Living Organization

The transition from static taxonomies to generative ontologies is not merely a technical upgrade; it is a philosophical shift. It is an admission that in the era of Generative AI, an organization is not a collection of fixed roles, but a living, breathing ecosystem of capabilities.

Leaders who continue to manage via "The Matrix" will find themselves captaining a ghost ship perfectly documented, but immobile. Those who embrace the ontology will possess the agility to reconfigure their entire workforce at the speed of innovation. The matrix is dead. Long live the graph.

GET YOUR SUBSCRIPTION TODAY!



NEVER MISS AN ISSUE!

- Enjoy huge savings
- Free home delivery
- Get your copy before everyone else

VISIT US ONLINE

- Catch the freshest features
- Updated daily
- Read anytime, anywhere

Join SkillPulse Community





SKILLPULSE

EXCLUSIVE MAGAZINE Data & AI Leaders

MAY 2026

ISSUE NO.01