
Using Agentic AI for Developing Engineers

The Role of Self-Healing AI Ecosystems in Technical Learning

David Swanagon

Article Info

Keywords

Machine Leadership
AI Agents
AI for HR
Domain Adapted LLMs
Competency Management
RAG-A
OTJ Task Books
Engineering Talent

Abstract

Abstract

This paper details the implementation of an AI Agent and MLOps architecture designed to modernize Competency Maps (C-MAPs), On-the-Job (OTJ) Task Books, and technical skills assessments within an engineering organization. In data-driven organizations, static learning documents often lead to progressive knowledge decay. This solution utilizes a Lang Graph-driven Retrieval-Augmented Generation (RAG-A) framework to create a "self-healing" talent ecosystem. By autonomously monitoring industry standards (Professional Societies, Tier 1 Competitors, Patent Filings) and internal SOPs, the Agent proactively suggests competency updates and personalized learning paths based on the domain.

The architecture is designed for multi-HRIS compatibility, providing seamless API orchestration with SAP, Oracle, and Workday. This ensures that talent intelligence is embedded directly into existing workflows. Supported by a robust MLOps pipeline—including model versioning via MLFlow, GitHub actions, and automated evaluation—the system maintains high accuracy and security within private cloud environments. By integrating Human-in-the-Loop (HITL) governance, the AI serves as a high-fidelity "Co-Pilot" for HR and Engineering leaders. This transition from manual learning updates to proactive talent intelligence significantly reduces time-to-autonomy, enhances process reliability, and secures the technical succession pipeline for data-driven organizations.

© 2025 David Swanagon. All Rights Reserved.

Introduction

In the age of AI, technical knowledge is constantly decaying at an unprecedented rate. STEM reliant organizations such as plant refineries, cybersecurity organizations, pharmaceutical researchers and so forth often find that their industrial hardware remains operational for decades, while their technical talent continues to struggle to keep up with innovation trends. For Chief Human Resources Officers (CHROs), the traditional manual approach to competency management—relying on static spreadsheets, SME interviews, and LMS updates—has become a significant operational risk. Research suggests that in safety-critical sectors, the "half-life" of technical knowledge is now approximately 2.5 to 5 years, meaning that without a digital trigger, up to 80% of technical standards can become obsolete or misaligned with field operations within just 24 months (Arbesman, 2012; ASEE, 2022).

This paper introduces a transformative solution: The Agentic Talent Intelligence System. Unlike traditional automation, which simply executes pre-defined tasks, Agentic AI possesses "agency"—the ability to reason, plan, and act autonomously to achieve complex goals (McKinsey, 2025; PwC, 2025). By shifting from "Talent on Paper" to a live, digital ecosystem, organizations can bridge the gap between "Work-as-Imagined" in a technical handbook and "Work-as-Done" on the refinery floor (Deloitte, 2023). This is essential as organizations that are exposed to safety risks or cybersecurity incidents must have processes that continuously align with market trends.

The introduction of AI Agents represents a structural shift for the HR function. HR Business Partners and Talent CoEs are no longer transactional administrators; they are becoming Strategic Talent Architects (SAP, 2025). By leveraging MLOps, the same discipline that is used for managing physical machinery reliability can be deployed for AI. This ensures that Competency Maps (C-MAPs), Technical Assessments, and On-the-Job (OTJ) Task Books are "Self-Healing." When a global engineering standard (e.g., ASME, NIST) is updated, the Agent proactively identifies the delta, drafts a revised competency program, and alerts the relevant supervisor (IBM, 2025; SHRM, 2021).

This paper introduces a robust technical methodology for building an "Agentic" talent academy. This includes exploring the integration of high-fidelity field assessments, SAP/Oracle/Workday synchronization, and the secure, "Air-Gapped" MLOps architecture required to protect sensitive employee data. The goal of this paper is to demonstrate how the partnership between HR and IT can eliminate knowledge decay, accelerate time-to-autonomy for new talent, and secure the technical succession pipeline in perpetuity (Gartner, 2023; World Economic Forum, 2025). Additional lines of inquiry are noted including transitioning from "Job-Based" to "Skill-Based" architectures and using Distributed Edge AI and wearables to automatically assess and update OTJ (On-The-Job) task books without human supervision.

Theoretical Framework

The design of the Agentic Academy is grounded in Sociotechnical Systems Theory (STS), which stipulates that organizational performance is optimized only when the social (human talent) and technical (AI systems) components are designed to work in synergy (Trist & Bamforth, 1951). In data-driven organizations, this synergy is maintained through dynamic capabilities, an

organization's ability to integrate, build, and reconfigure internal and external competences to address rapidly changing environments (Teece et al., 1997). By utilizing Agentic AI, the organization moves from a Static to Adaptive Competency model. Here, the system acts as a digital nervous system, sensing shifts in global standards and automatically recalibrating the workforce's "Zone of Proximal Development"—the distance between what a professional can do alone and what they can do with expert (or AI) guidance (Vygotsky, 1978).

To describe the theoretical framework, this paper imagines a theoretical use case focusing on a Cybersecurity Company that specializes in critical infrastructure protection. The challenge the company is facing is a recurring knowledge decay that happens within 18 months. The company employs 1,500 Security Operations Center (SOC) analysts. In the cybersecurity domain, the average lifespan of a technical skill is estimated at just 18 to 24 months before it is superseded by new threat vectors (World Economic Forum, 2023). Before implementing Agentic AI, the company relied on manual quarterly audits to update their Competency Maps (C-MAPs), resulting in a "Readiness Gap" where analysts were often trained on vulnerabilities that had already been patched or replaced by new "Zero-Day" exploits. This made the OTJ books meaningless.

To address this problem, the company built an AI Agent using the framework described in the methodology section and integrated this model with their Workday Skills Cloud. The process included a structured sensing, mapping, and activation layer. First, the agent would continuously monitor the NIST vulnerability database and MITRE attack frameworks. As new tactics are introduced for critical threats, the Agent autonomously identifies that only 14% of SOC analysts have maintained the log analysis competencies to detect the new threat. (Deloitte, 2024). The AI Agent drafts a new OTJ Task Book entry and pushes the micro-learning to all affected engineers. The HR Professional serves as the Human-In-the-Loop to approve learning changes.

The output of this theoretical example is that the Time-To-Train new SOC analysts will decrease. Furthermore, the speed at which technical updates (i.e., NIST, MITRE) are integrated into the learning framework and day-to-day operations will increase. This aligns with industry data showing that organizations using AI-driven skill mapping are 1.5x more likely to report high organizational agility (Deloitte, 2024).

This theoretical use case serves as a blueprint for the "Agentic Academy," demonstrating that technical development in data-driven environments must move from a periodic manual review to adaptive learning. By embedding LangGraph-driven reasoning directly into the competency lifecycle, the company successfully transitioned from a reactive training posture to a proactive defense posture, effectively neutralizing the 18-month skill obsolescence cycle (World Economic Forum, 2023). This outcome underscores the fundamental shift in the CHRO-CIO relationship. When talent intelligence is synchronized with real-time industry benchmarks, the workforce becomes a self-healing asset capable of navigating the "Agentic Divide" (Deloitte, 2024). Ultimately, this use case conceptually proves that the integration of Agentic AI is not merely a technical upgrade, but a strategic imperative for any organization where the speed of learning must exceed the speed of change to ensure institutional resilience and operational safety. That said, more data is needed on real-world applications to support these findings in practice.

Methodology

Building an Agentic AI Talent Ecosystem

The AI Agentic Academy provides a structural reinforcement for Cognitive Load Theory (CLT) and the 70-20-10 Model of professional development. By utilizing AI to identify real-time skill gaps, the system ensures that learning occurs within the Zone of Proximal Development, providing engineers with "Just-in-Time" information rather than "Just-in-Case" training. This indicates AI agents should not be viewed as novelty tools detached from established learning science. Instead, they support best practices that have long been established for developing technical teams. For example, the AI Agent engages in scaffolding and micro-learning activities by decomposing complex C-MAP requirements into manageable OTJ (On-The-Job) task books. This ensures that engineers are not overloaded with materials, while receiving the targeted tasks that drive development. Separately, the AI Agent's feedback loop provides immediate responses to learning challenges. This helps engineers iterate through new material, especially scenario modeling. This is a core requirement for transitioning from novice to expert in high-hazard engineering environments (Ericsson & Pool, 2016). Another benefit is spaced repetitive and retrieval practice. The MLOps "Maintenance Loop" identifies when an engineer has not performed a specific task in 12 months, proactively triggering a "Refresher Task" to combat the Forgetting Curve (Ebbinghaus, 1885).

We begin by reviewing the multi-layered technical design requirements for implementing an "Agentic" Talent Academy. This is followed by a review of the learning best practices that support the system's architecture. This methodology moves beyond traditional software deployment into Agentic Engineering, where the system is designed to perceive changes in the industrial and cybersecurity environment and execute complex reasoning loops without constant human prompting, oversight, or decision-making. This methodology moves beyond traditional software deployment into agentic engineering, where systems perceive environmental changes and execute complex reasoning loops without constant human prompting or oversight (McKinsey & Company, 2025).

Step 1: Data Ingestion

The foundational layer of this system involves the aggregation and "vectorization" of engineering knowledge. Unlike legacy systems that store data in flat files, the data structures supporting an AI Talent Academy creates a high-dimensional mathematical map of all relevant technical standards, coding repositories, safety guidelines, and behavioral traits. This combats the "half-life" of engineering facts (Arbesman, 2013). The process requires ingesting three data sources including external regulatory guidelines (i.e., ASME, NIST), internal standard operating procedures (SOPs), and historical Competency Maps (C-MAPs) for each technical domain.

Step 2: Vector Embedding

From this, a vector embedding model such as text-embedding 3 is used to convert unstructured text into vectors that are stored in a database (Qdrant, 2023). This enables the agent to utilize Retrieval-Augmented Generation (RAG) to provide grounded, non-hallucinatory answers based on specific source documents (Pinecone, 2023). These tools are particularly useful

as they help create conceptual understanding of data without transmitting sensitive data to the public cloud. Essentially, vector embeddings help the AI agent understand that a prompt written in general terms is still contextually connected to a specific regulatory guideline.

Step 3: LangGraph

The next step is developing the reasoning engine. LangGraph is an effective framework as it allows the agent to remember previous progress, use loop back procedures for corrections, and handle cycles that do not conform to linear AI principles (LangChain, 2023). For example, the AI Agent may draft an OTJ task book for a cybersecurity engineer, which is rejected by the assigned technical SME. LangGraph allows the agent to review the feedback, loop back to its design phase, and develop a revised draft for secondary review. This feature allows for Human-in-the-Loop (HITL) design principles to be embedded in key decision points (Fails & Olsen, 2003). This means the AI Agent cannot update a competency map, mentor-mentee match recommendation, or technical assessment report without obtaining a digital signature from a human operator (i.e. HR Business Partner).

Step 4: Domain Adapted LLM and LoRA

Once vectorization and the LangGraph framework has been established, the next step is to pursue a Domain-Adapted LLM model. This process optimizes base language models such as GPT-4, Claude, and others that understand broad engineering principles, while lacking specific details of an organization's technology stack. This process is typically executed using LoRA (Low-Rank Adaptation) to fine-tune the base model by offering 750 – 1,000 technical scenarios that help the AI model learn the specific details of the organization, culture, and data architecture (Microsoft, 2023). This ensures the AI model understands the contextual priorities, jargon, and organizational principles, which ensures responses are accurate when applied to the company's unique tasks.

Step 5: Integrate AI Agent with HRIS

The next step involves integrating the AI Agent with the HRIS system through API orchestration. A token is provided to the AI Agent granting "read/write" access to the system of record. For example, an OData API would be used for SAP, while Rest APIs and Workday Web Services would be used for Oracle and Workday, respectively. This process allows the AI Agent to synchronize job profiles, learning management systems, and skills clouds with external market trends and other internal sources such as SOPs and data security guidelines. The agent's value is realized through its ability to synchronize with the "System of Record." The methodology adapts to the organization's specific HRIS infrastructure (Gartner, 2023).

Step 6: MLOPs and Maintenance Loops

The final step in the process is the MLOPs and Maintenance loop. This process ensures that "Data Drift" and degradation does not occur because of the API losing accuracy (IBM, 2025). Tools such as MLFlow for experiment tracking, RAGAS for assessing RAG relevance and accuracy, and Evidently AI ensure that technical teams can quickly detect and address performance

declines (Evidently AI, 2023). When a trigger is issued, the AI Agent retrains the LoRA domain-adapted LLM process, which ensures the prompt responses remain aligned with market trends.

Literature Review

The evolution of talent management has reached a critical juncture where manual, static competency frameworks are insufficient for the demands of data-driven organizations such as Oil & Gas, Pharmaceuticals, and Cybersecurity. Current literature emphasizes the importance of shifting toward dynamic, adaptive learning systems that utilize artificial intelligence to mitigate knowledge decay and enhance organizational agility.

Technical Knowledge Decay

A core challenge in engineering and cybersecurity industries is the accelerating rate of information obsolescence. Arbesman (2013) introduced the concept of the "half-life of facts," arguing that technical knowledge in STEM fields has a measurable expiration date, requiring constant institutional recalibration. This is supported by the American Society for Engineering Education (ASEE, 2022), which stipulates that without continuous digital reinforcement, a significant portion of engineering standards become misaligned with field reality within two years.

Agentic AI and Autonomous Reasoning

Advancements in Large Language Models (LLMs) have shifted the focus from generative AI to goal-oriented reasoning. Unlike traditional automation, agents utilize cyclic logic and reasoning loops to manage complex, multi-step tasks such as curriculum design and skill mapping (LangChain, 2023). McKinsey & Company (2025) reports that "Agentic" workflows—those capable of self-correction and external tool use—represent the next frontier in operational efficiency, allowing systems to act as "Reasoning Engines" rather than mere data repositories.

Sociotechnical Synergy and Human-in-the-Loop (HITL)

The integration of AI into talent development requires a sociotechnical lens. This means balancing technical efficiency with human oversight (Trist & Bamforth, 1951). Fails and Olsen (2003) highlighted the importance of "Interactive Machine Learning," a forerunner to modern Human-in-the-Loop (HITL) frameworks. These frameworks ensure that while AI can draft competency models (C-MAPs), Subject Matter Experts (SMEs) retain the final "Approval" authority, which preserves institutional safety and accountability (IBM, 2025).

Spaced Repetition and Cognitive Scaffolding

Modern learning science supports the use of AI for "Just-in-Time" competency verification. Ericsson and Pool (2016) emphasize that spaced retrieval practice and immediate feedback are the primary drivers of expertise acquisition. AI agents facilitate this by delivering micro-learning modules precisely when an engineer encounters a skill gap. Furthermore, the application of Ebbinghaus' (1885) forgetting curve suggests that automated "Refresher Tasks" are

essential for long-term retention in high-stakes environments, a process now digitized through MLOps-driven maintenance loops (Murre & Dros, 2015).

Skills-Based Organization

The literature from Tier 1 consultancies indicates a structural shift in the HR function itself. Deloitte (2024) argues that the "Skills-Based Organization" is the new operating model for the future, where AI agents map internal capabilities to market demands in real-time. This transition requires a robust MLOps (Machine Learning Operations) architecture to ensure that the "Talent Intelligence" remains accurate, secure, and integrated with core HRIS systems like SAP or Workday (Gartner, 2023; World Economic Forum, 2025)

Summary of Findings

Using the fictitious use case example, the integration of Agentic AI into the engineering competency lifecycle represents more than a software upgrade. It is a strategic solution to the problem of Technical Knowledge Decay. By utilizing LangGraph-driven reasoning and MLOps-monitored learning loops, organizations can ensure their workforce remains perpetually aligned with global benchmarks like ASME, NIST, and other technical trends. To transition from a static, manually driven talent framework to an Agentic Academy, leadership must synchronize technical deployment with organizational change management. Based on the frameworks established in this paper, the following recommendations provide a roadmap for CHROs and CIOs.

Establish a "Human-in-the-Loop" (HITL) Governance Board. AI should never operate with complete autonomy, particularly in safety-critical engineering environments. HR and IT organizations should form cross-functional talent councils, comprising SMEs, HR CoE leaders, and IT Security Officers. This council can define the goals of the Agentic Academy, while serving as the final decision-making for C-MAP, OTJ, and micro-learning updates.

Adopt a "Skills-First" Data Architecture. Legacy HRIS systems are often structured around static job titles which mask the underlying technical capabilities of the workforce. Furthermore, the job titles do not evolve and adapt at the same pace of skill requirements. To stay aligned with market trends, organizations should engage in skills mapping within the HRIS system. This includes utilizing tools such as Workday Skills Cloud or SAP SuccessFactors Talent Hub.

Invest in "Air-Gapped" MLOps Infrastructure. Data protection and IP security is paramount for protecting proprietary engineering SOPs and employee performance data. Organizations should deploy the AI Agent and its Vector Database within a private cloud environment. This ensures that the domain adapted training activities are not used for public LLM training sets.

Pilot "Just-in-Time" OTJ Verification. Transitioning from annual reviews to continuous assessment requires a shift in how field performance is captured. CHROs can pilot this concept by focusing on critical technical domains such as Process Engineering and Threat Intelligence. This includes using the AI Agent to push micro-learning tasks to supervisor's mobile dashboards. By doing this, organizations can validate Time-To-Train improvements before scaling to other areas.

Reimagine the HR-IT Partnership. The traditional SLA (help desk ticket) relationship between HR and IT is insufficient for agentic systems. Moving forward, CHROs and CIOs should create joint intelligence teams that define the logic of the domain adapted LLM systems, while ensuring the data pipelines, MLOps, and governance activities are managed effectively.

Conclusions

The transition to an Agentic Academy represents a fundamental shift from manual, reactive talent management to a proactive model of adaptive learning. By linking the competency lifecycle to a robust MLOps architecture and LangGraph-driven reasoning, engineering organizations can finally overcome the persistent challenge of technical knowledge decay (Arbesman, 2013). This architecture creates a "self-healing" workforce that evolves in line with global engineering standards and real-time operational demands. As the "Agentic Divide" continues to widen, the ability to synchronize human expertise with autonomous talent intelligence will define the next generation of industrial leaders (McKinsey, 2025). This requires reimagining the CHRO and CIO relationship in a way that fosters joint intelligence gathering. By engaging in these activities, organizations can secure their technical succession pipelines, minimize risks, and ensure that their workforce is prepared for the future.

Additional Lines of Inquiry

This paper recommends further investigation of several research questions. These inquiries should focus on the intersection of human psychology, machine learning stability, and organizational economics. This will ensure that solutions scale in an ethical manner. First, companies must assess the psychological impact of continuous assessment. Traditional competency models rely on periodic evaluations, whereas agentic systems enable a state of continuous review. Further research is required to determine how AI-driven OTJ tasks affect employee psychological safety and long-term engagement (Ericsson & Pool, 2016). Separately, it is important that standard HR metrics evolve to quantify knowledge. Lastly, the issue of Agent-to-Agent collaboration is important as the AI ecosystem continues to expand. This means evaluating how an AI Talent Agent interacts with an Operational Agent.

References

- Arbesman, S. (2013). *The half-life of facts: Why everything we know has an expiration date*. Current/Penguin Publishing Group.
- American Society for Engineering Education (ASEE). (2022). *The evolving nature of engineering work: Benchmarking technical knowledge decay*. Washington, DC: ASEE Press.
- Deloitte. (2024). *The skills-based organization: A new operating model for work and the workforce*. Deloitte Insights. <https://www2.deloitte.com/us/en/insights/topics/talent/special-edition-skills-based-organization.html>
- Deloitte. (2023). *2023 engineering and construction industry outlook*. Deloitte Insights. <https://www.deloitte.com/global/en/Industries/industrial-construction/analysis/gx-engineering-and-construction-industry-trends.html>
- Ebbinghaus, H. (1885). *Memory: A contribution to experimental psychology* (H. A. Ruger & C. E. Bussenius, Trans.). Teachers College, Columbia University. (Original work published 1885).
- Ericsson, A., & Pool, R. (2016). *Peak: Secrets from the new science of expertise*. Houghton Mifflin Harcourt.
- Evidently AI. (2023). *ML observability: Monitoring data and model drift*. <https://evidentlyai.com>
- Fails, J. A., & Olsen, D. R. (2003). Interactive machine learning. *Proceedings of the 8th International Conference on Intelligent User Interfaces*, 39–46.
- Gartner. (2023, October 18). *Magic quadrant for cloud HCM suites for 1,000+ employee enterprises* (R. Chandra & S. Grinter, Eds.). Gartner, Inc.
- IBM. (2025). *AI academy: Prioritizing AI investments that drive business growth*. IBM Think Insights. <https://www.ibm.com/think/insights/ai-skills-you-need-for-2025>
- LangChain. (2023). *LangGraph: Building stateful, multi-actor applications with LLMs*. LangChain Blog. <https://blog.langchain.dev/langgraph/>
- McKinsey & Company. (2025, September 12). *One year of agentic AI: Six lessons from the people doing the work* (L. Yee, M. Chui, R. Roberts, & S. Xu, Eds.). McKinsey QuantumBlack. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/one-year-of-agentic-ai-six-lessons-from-the-people-doing-the-work>
- Microsoft. (2023). *LoRA: Low-rank adaptation of large language models*. Microsoft Research.
- MLflow. (2023). *MLflow: A platform for the machine learning lifecycle*. <https://mlflow.org/docs/latest/index.html>
- Murre, J. M., & Dros, J. (2015). Replication and analysis of Ebbinghaus' forgetting curve. *PLOS ONE*, 10(7), e0131340.
- Pinecone. (2023). *Vector databases for the enterprise: A technical overview*. Pinecone Systems.
- Pinecone. (2023). *Retrieval-Augmented Generation (RAG) explained*. <https://www.pinecone.io/learn/retrieval-augmented-generation/>
- PwC. (2025). *PwC's 2025 US responsible AI survey: From policy to practice*.
- PricewaterhouseCoopers. <https://www.pwc.com/us/en/tech-effect/ai-analytics/responsible-ai-survey.html>
- Qdrant. (2023). *Vector database for the enterprise*. <https://qdrant.tech>
- SAP. (2025). *Looking back and moving forward: HR trends from 2024 to 2025*. SAP SuccessFactors Research. <https://www.sap.com/research/hr-trends-from-2024>
- Society for Human Resource Management (SHRM). (2021). *SHRM body of applied skills and knowledge (BASK)*. SHRM. <https://www.shrm.org/credentials/certification/exam-preparation/bask>

- Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509–533. [https://doi.org/10.1002/\(SICI\)1097-0266\(199708\)18:7<509::AID-SMJ882>3.0.CO;2-Z](https://doi.org/10.1002/(SICI)1097-0266(199708)18:7<509::AID-SMJ882>3.0.CO;2-Z)
- Trist, E. L., & Bamforth, K. W. (1951). Some social and psychological consequences of the longwall method of coal-getting. *Human Relations*, 4(1), 3–38. <https://doi.org/10.1177/001872675100400101>
- Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Harvard University Press.
- Workday. (2023). *Workday API reference guide*. <https://workday.com>
- World Economic Forum. (2023). *The future of jobs report 2023*. <https://www.weforum.org/publications/the-future-of-jobs-report-2023/>
- World Economic Forum (WEF). (2025, January 7). *The future of jobs report 2025*. World Economic Forum. <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>