



*An Online PDH Course
brought to you by
CEDengineering.com*

Implementing Agentic AI in Engineering

Course No: W03-001

Credit: 3 PDH

Brian Lisiewski, P.E.



Continuing Education and Development, Inc.

P: (877) 322-5800

info@cedengineering.com

www.cedengineering.com

Table of Contents

1. Introduction to AI and Agentic AI	1
1.1 What is Agentic AI?.....	3
2. Conceiving and Designing AI Agents	6
2.1 Identifying AI Opportunities (Use Case Conception)	6
2.2 Defining Goals and Requirements	8
2.3 AI Agent Design Principles	9
3. AI Risk Management and Security Planning	11
3.1 Identifying Risks of AI in Engineering.....	11
3.2 Risk Mitigation Strategies and Safe Design	13
4. Implementing AI in Engineering Design Processes	16
4.1 Pilot Projects and Gradual Integration.....	16
4.2 Integrating AI into the Engineering Workflow.....	17
4.3 Case Study Examples.....	20
5. Developing KPIs for AI Performance and Impact	22
5.1 Why KPIs Matter for AI Projects	22
5.2 Categories of KPIs for AI	23
5.3 Monitoring and Continuous Improvement.....	24
6. Policy and Governance for AI Implementation	25
6.1 Developing an AI Policy.....	25
6.2 Governance Structure and Roles.....	27
7. Developing a Safe AI Implementation Action Plan	29
8. Conclusion and Next Steps	30

Table of Figures

Figure 1: Historical Timeline of Artificial Intelligence Development (1956–2020s) 1

Figure 2: Traditional vs. Agentic AI..... 5

Figure 3: AI Project Conception: Structured Planning Workflow..... 6

Figure 4: AI Risk Management..... 13

Figure 5: AI Integration in Engineering Workflow 17

Figure 6: Key Performance Indicators (KPIs) for Evaluating AI Performance and Impact..... 22

Figure 7: AI Governance Structure: Roles and Organizational Oversight 27

Figure 8: Safe AI Implementation Plan 29

1. Introduction to AI and Agentic AI

Artificial Intelligence (AI) refers to a broad field of computer science focused on creating systems capable of performing tasks that typically require human intelligence, such as **learning from data, making decisions, recognizing patterns, and solving problems**^{1 2}. Modern AI encompasses various subfields: *machine learning* (algorithms that learn from data), *natural language processing* (understanding/generating human language)³, *computer vision* (interpreting images and video)⁴, *robotics*, and more. AI has evolved through several eras, from early rule-based **expert systems** to today’s data-driven **deep learning** and **generative AI** models.

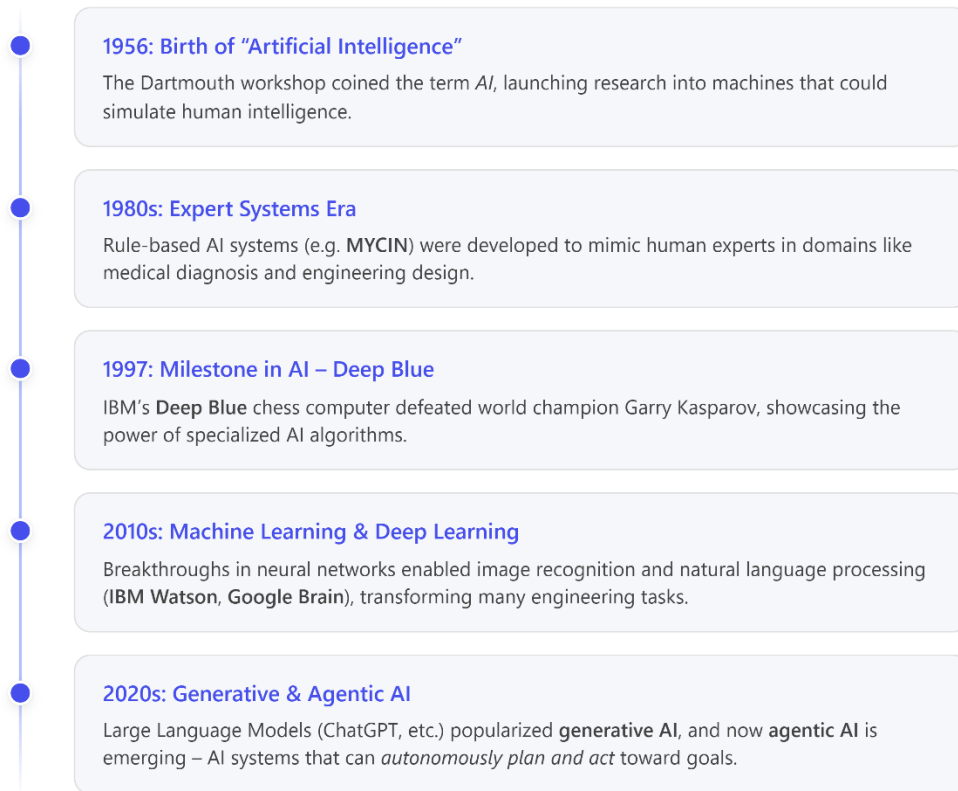


Figure 1: Historical Timeline of Artificial Intelligence Development (1956–2020s)

¹https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf
²https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf
³https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf
⁴https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

Relevance to Engineering: AI has become a powerful enabler in nearly every engineering discipline. By leveraging AI, engineers can **analyze large data sets for insights, automate repetitive tasks, optimize designs, and predict outcomes** with greater accuracy^{5 6}. For example, contemporary engineering applications of AI include:

- *Generative design software* that **iterates thousands of design options** (e.g. building layouts or component designs) under given constraints, helping engineers explore innovative solutions rapidly⁷.
- *Predictive analytics and simulations* that improve **project scheduling and risk management** by forecasting delays or cost overruns and suggesting optimized schedules⁸.
- *Computer vision systems* for **automated safety and quality inspections**, like using drone imagery to detect construction defects or ensure workers wear proper safety equipment⁹.
- *Predictive maintenance in infrastructure*, where AI models analyze sensor data (vibration, temperature, etc.) to predict equipment failures in advance, reducing downtime¹⁰.

Engineering firms are embracing AI not only for efficiency and cost reduction, but also to enhance **design quality, safety, and innovation**, thereby gaining competitive advantage¹¹. **However, adopting AI must be done responsibly**, without proper controls, AI mistakes or biases could lead to errors or safety incidents that undermine trust and public safety¹². This need for caution and governance is especially true for the new frontier of *agentic AI*.

⁵https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁶https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁷https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁸https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁹https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

¹⁰https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

¹¹https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

¹²https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

1.1 What is Agentic AI?

Agentic AI is an advanced form of AI characterized by **autonomous decision-making and goal-directed behavior**^{13 14}. Traditional AI systems (including many **generative AI** models) typically operate on an input, output basis: they respond to user queries or perform narrowly defined tasks under close human supervision¹⁵. In contrast, an *agentic AI system*, often composed of multiple interacting **AI agents**, can **independently plan its actions, make decisions, and execute tasks** to achieve a specified goal with minimal human intervention^{16 17}. The term “agentic” highlights that these AI have **agency**: the capacity to **act purposefully** on their own¹⁸.

In practical terms, an agentic AI behaves more like a **proactive collaborator** than a passive tool. It can break down a high-level goal into subtasks, coordinate various AI components or sub-agents, and carry out multi-step plans while dynamically adapting to new information^{19 20}. For example, rather than simply analyzing data when asked, an agentic AI might detect an anomaly in a structural health monitoring system, **decide** to gather additional sensor data and weather records, run simulations to diagnose the issue, and then raise an alert with recommended actions, all **without needing a new prompt for each step**. It moves “**from helping to doing**”²¹.

Key characteristics of agentic AI include:

- **Autonomy and Proactivity:** They operate **without constant human oversight**, maintaining and pursuing long-term goals, initiating actions as needed^{22 23}. Instead of waiting for instructions at each step, agentic AI can decide **what needs to be done next** to meet its objectives²⁴.
- **Adaptability:** Agentic AI systems can **learn from real-time feedback and past experiences**, adjusting their strategies to improve over time^{25 26}. They often incorporate *feedback loops* to refine their decisions and performance continually.
- **Collaborative Behavior:** They may consist of **multiple specialized agents working together** (a *multi-agent system*) orchestrated towards a common goal²⁷. An agentic AI can also interact with humans and other systems, requesting help or data when needed²⁸.

¹³<https://www.ibm.com/think/topics/agentic-ai>

¹⁴<https://www.oracle.com/artificial-intelligence/agentic-ai/>

¹⁵<https://www.oracle.com/artificial-intelligence/agentic-ai/>

¹⁶<https://www.ibm.com/think/topics/agentic-ai>

¹⁷<https://www.oracle.com/artificial-intelligence/agentic-ai/>

¹⁸<https://www.ibm.com/think/topics/agentic-ai>

¹⁹<https://www.ibm.com/think/topics/agentic-ai>

²⁰<https://www.oracle.com/artificial-intelligence/agentic-ai/>

²¹<https://www.oracle.com/artificial-intelligence/agentic-ai/>

²²<https://www.ibm.com/think/topics/agentic-ai>

²³<https://www.ibm.com/think/topics/agentic-ai>

²⁴<https://www.oracle.com/artificial-intelligence/agentic-ai/>

²⁵<https://www.ibm.com/think/topics/agentic-ai>

²⁶<https://www.oracle.com/artificial-intelligence/agentic-ai/>

²⁷<https://www.ibm.com/think/topics/agentic-ai>

²⁸<https://www.oracle.com/artificial-intelligence/agentic-ai/>

- **Goal-Driven Reasoning:** Instead of focusing on one fixed task, agentic AI uses techniques (like planning algorithms or reinforcement learning) to **choose among many possible actions** to best achieve the overall objective^{29 30}. This is akin to a project manager allocating tasks to team members and adjusting plans on the fly, rather than a single tool performing one function^{31 32}.

Why Now? Agentic AI is emerging now thanks to recent breakthroughs in AI: notably the rise of **Large Language Models (LLMs)** and other advanced AI components which can be combined as “brains” and “tools” for agents^{33 34}. For instance, an LLM-based “conductor” agent can oversee strategy and communication, while specialized sub-agents handle tasks like data retrieval, analysis, or design generation³⁵. This builds on generative AI’s ability to create content by adding the ability to **take actions based on that content**^{36 37}. As a result, today’s AI has “leaped” from merely providing recommendations to actually **executing complex sequences of actions** toward a goal³⁸.

Engineering Example: *Imagine an engineering project management agentic AI* that monitors project progress. It could automatically track task completion against the schedule, identify delays, diagnose causes (e.g. resource shortages or technical blockers), and then autonomously reallocate resources or adjust timelines in the project management system to keep the project on track , notifying human managers only as needed. Such an agent could save significant time and reduce errors by proactively handling routine decisions, allowing engineers to focus on critical creative and safety-critical tasks.

²⁹<https://www.ibm.com/think/topics/agentic-ai>

³⁰<https://www.ibm.com/think/topics/agentic-ai>

³¹<https://www.oracle.com/artificial-intelligence/agentic-ai/>

³²<https://www.oracle.com/artificial-intelligence/agentic-ai/>

³³<https://www.ibm.com/think/topics/agentic-ai>

³⁴<https://www.ibm.com/think/topics/agentic-ai>

³⁵<https://www.ibm.com/think/topics/agentic-ai>

³⁶<https://www.ibm.com/think/topics/agentic-ai>

³⁷<https://www.ibm.com/think/topics/agentic-ai>

³⁸<https://www.oracle.com/artificial-intelligence/agentic-ai/>

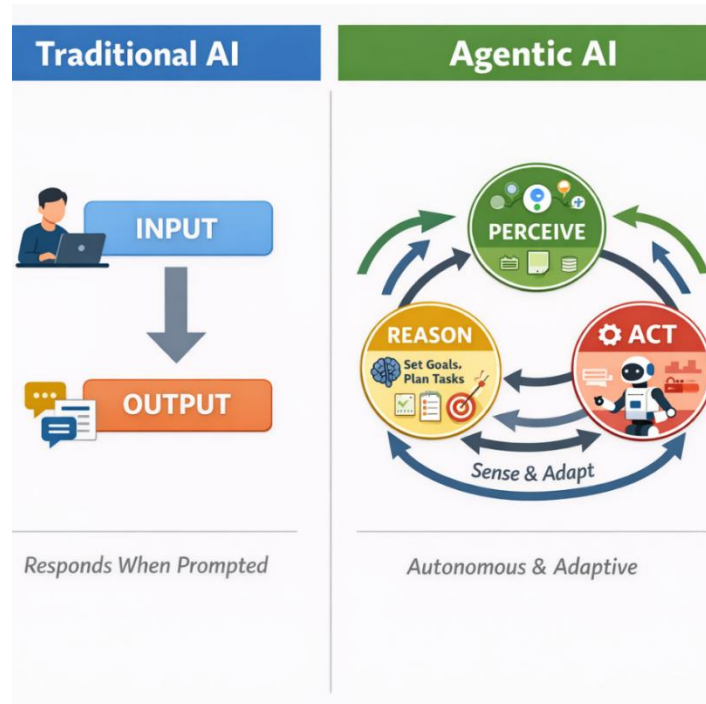


Figure 2: Traditional vs. Agentic AI

Caution: While agentic AI offers powerful capabilities, its autonomy also poses **new challenges**, especially in safety-critical fields like engineering. Without proper constraints, an AI agent might take inappropriate actions or make unsupervised decisions that conflict with regulatory requirements or ethical norms. **Therefore, implementing agentic AI in an engineering firm requires a disciplined, safety-focused approach.** The rest of this course will outline a step-by-step process, from initial conception and design through risk management, deployment, performance monitoring, and governance, to harness the benefits of agentic AI while controlling its risks.

2. Conceiving and Designing AI Agents

Implementing agentic AI in an engineering company **starts with careful planning and design**. In this module, we cover how to **identify suitable opportunities for AI (“conceiving” an AI solution) and design an AI agent or system** to meet your engineering objectives. Key steps include **selecting high-impact use cases, defining the AI’s goals and scope, choosing the right methodologies/technologies, and architecting the agent’s structure**.

2.1 Identifying AI Opportunities (Use Case Conception)

Before building any AI solution, it’s crucial to **conceive the right problem for the AI to solve**. Engineers should ask: *Where can an AI agent provide the most value in our processes?* Look for tasks that are **repetitive, data-intensive, or prone to human error**, as well as problems where improved prediction or optimization could yield big benefits^{39 40}. **High-impact, low-risk use cases** make ideal starting points⁴¹. For example:

- **Design optimization tasks:** e.g. using AI to generate and evaluate many design alternatives (in architecture, mechanical part design, circuit layouts, etc.), which can dramatically speed up early-stage design exploration.
- **Routine analysis or QA checks:** e.g. an AI that reviews engineering drawings or code for compliance with standards, flagging errors for human engineers to fix.
- **Monitoring and alerting:** e.g. an AI agent watching sensor data from infrastructure or manufacturing processes, which detects anomalies and triggers maintenance actions.



Figure 3: AI Project Conception: Structured Planning Workflow

³⁹https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁴⁰https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁴¹https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

Selecting a focused application with clear success metrics (schedule compliance, error reduction, cost savings, etc.) makes it easier to measure the AI's benefits and convince stakeholders of its value⁴². It also limits risk while your team is new to AI. **Start small with 1 , 2 pilot use cases** that can be tightly controlled and observed⁴³.

As you identify opportunities, secure **buy-in from leadership and relevant departments** early. Form a **cross-functional team**, drawing on engineering, IT/data, project management, and leadership, to champion the AI initiative⁴⁴. This team will guide the project, set objectives, and ensure the AI solution aligns with business goals. Early leadership support is critical to allocate resources and signal the importance of the project to the whole organization⁴⁵.

Insight: When defining an AI project, focus on *business goals over technology*. For example, rather than “we need a machine learning system,” frame it as “we want to reduce design cycle time by 20%” or “improve safety inspection coverage.” This keeps the effort grounded in solving real engineering problems, not adopting AI for AI's sake.

⁴²https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁴³https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁴⁴https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁴⁵https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

2.2 Defining Goals and Requirements

Once a promising use case is chosen, clearly **define the objective** of the AI agent in that context. What **specific goal** should the agent achieve or assist with? For instance: “*Optimize the structural design of a bridge for minimal material cost while meeting safety factors,*” or “*Monitor pipeline sensor data and flag likely leak events in real time.*” A well-defined goal guides the entire development process and evaluation of success⁴⁶.

Also establish the **requirements and constraints** for the AI system:

- **Scope & Boundaries:** Delineate what the agent will and will *not* do. This prevents project creep and unsafe autonomy. For example, an agent might be allowed to propose design changes but not approve them without human review, or it may issue maintenance recommendations but not directly shut down a production line. Setting **clear scope limits** is a key safety measure⁴⁷.
- **Performance Metrics:** Determine what metrics the AI must achieve (accuracy, speed, etc.) for the solution to be acceptable. For instance, a defect-detecting vision AI might need to catch $\geq 95\%$ of critical defects with $< 1\%$ false alarms. These targets inform the design and later serve as KPIs to evaluate performance.
- **Data and Integration Needs:** Identify what data the agent will require (historical data sets, real-time sensor feeds, design specifications, etc.) and what systems it must interact with (CAD software, databases, project management tools). Ensure the necessary data is available, of good quality, and accessible (address issues of data **completeness, quality, and bias** early)^{48 49}. Engaging your IT/data management team here is vital.
- **User Interface:** Plan how engineers or other users will interact with the agent. Will it be through a chat interface, a dashboard, or integrated into existing software? **User experience design** is important so that the AI’s outputs are understandable and actionable for humans.

Throughout this phase, keep the **end-user (engineers, managers, etc.) involved**. Their domain expertise is critical for setting realistic goals and constraints. This also helps ensure the AI will be embraced by users, not viewed as a “black box” threat. Good communication can alleviate fears of job displacement by emphasizing the AI as a tool to **augment engineers, not replace them**⁵⁰.

⁴⁶<https://www.oracle.com/artificial-intelligence/agentic-ai/>

⁴⁷<https://www.oracle.com/artificial-intelligence/agentic-ai/>

⁴⁸<https://www.multimodal.dev/post/ai-kpis>

⁴⁹<https://www.multimodal.dev/post/ai-kpis>

⁵⁰https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁵¹https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

2.3 AI Agent Design Principles

With objectives and requirements defined, the next step is to design the AI agent or agents. Here we introduce fundamental principles and methodologies for **AI agent design**:

- **Intelligent Agent Model:** In AI theory, an **intelligent agent** is something that perceives its **environment** through sensors and **acts** upon that environment through actuators (or outputs) in order to achieve goals^{52 53}. For example, an AI software agent might “sense” by reading data from a database (its input sensor) and “act” by sending alerts or writing design recommendations (its output actuator). Always define what data your agent will sense and what actions it is allowed to take.
- **Autonomy vs. Control:** Decide how autonomous the agent should be. A fully autonomous agent makes decisions with no human in the loop, whereas a semi-autonomous agent might defer to humans at key decision points. For safety, **high-stakes decisions should have human oversight or approval** (at least until the AI has proven its reliability)^{54 55}. Many successful implementations start with the AI in an advisory role (e.g., providing recommendations that humans confirm) and only later increase autonomy as trust in the AI grows.
- **Single Agent or Multi-Agent System:** Determine if the problem is best solved by one agent or a **team of specialized agents working together**. In a **multi-agent system**, different agents can handle different tasks (for instance, one agent analyzes data, another generates design options, and a top-level agent orchestrates their work)^{56 57}. Multi-agent architectures add complexity, but they can mirror the way engineering teams delegate tasks. If using multiple agents, design a clear **coordination mechanism**, e.g., a central orchestrator or a decentralized protocol, so that all agents remain aligned to the overall goal⁵⁸.
- **Methodologies and Frameworks:** Leverage established **AI design methodologies** to guide development. For example, *agent-oriented software engineering* methodologies like **Gaia** or **Prometheus** provide structured approaches to define agent roles, behaviors, and interactions. Similarly, frameworks such as **OpenAI’s AutoGPT** or **LangChain** (for LLM-based agents) or **JADE** (Java Agent Development Framework for multi-agent systems) can accelerate development of agentic systems by providing pre-built components. The choice of tools should align with the team’s expertise and the problem needs.

⁵²<https://www.ibm.com/think/topics/agentic-ai>

⁵³<https://www.ibm.com/think/topics/agentic-ai>

⁵⁴<https://www.oracle.com/artificial-intelligence/agentic-ai/>

⁵⁵<https://www.oracle.com/artificial-intelligence/agentic-ai/>

⁵⁶<https://www.ibm.com/think/topics/agentic-ai>

⁵⁷<https://www.ibm.com/think/topics/agentic-ai>

⁵⁸<https://www.ibm.com/think/topics/agentic-ai>

- **Goal Decomposition and Planning:** If the agent’s task is complex, design how it will break the goal into sub-tasks and in what order to execute them. This often involves integrating **planning algorithms** or using LLMs to dynamically generate task lists⁵⁹. Ensure the agent can update its plan if conditions change or if some approach fails, this adaptability is a key advantage of agentic AI. Methods like **reinforcement learning** or planning with decision trees can enable agents to discover optimal actions to meet goals⁶⁰.
- **Knowledge and Learning:** Decide how the agent will **learn and store knowledge**. Will it use machine learning models trained on historical data? Will it have a knowledge base or rules (e.g., engineering codes/standards) it can reference? Many agentic AIs use a combination: for example, an agent might call a pre-trained predictive model for fast decisions and an LLM for reasoning or advice. Incorporating a feedback mechanism (e.g., the agent learns from each project to improve next time) can continuously enhance performance^{61 62}.

Throughout design, constantly consider **reliability and edge cases**. Ask “*What could go wrong?*” For every function the agent performs, think of failure modes: How will it handle missing or bad data? What if a sensor or input is unavailable? If the agent encounters an unexpected scenario, does it fail safely (e.g., alert a human) rather than take reckless action? Proactively designing for these questions is part of risk management, which we cover next.

⁵⁹<https://www.oracle.com/artificial-intelligence/agentic-ai/>

⁶⁰<https://www.ibm.com/think/topics/agentic-ai>

⁶¹<https://www.ibm.com/think/topics/agentic-ai>

⁶²<https://www.oracle.com/artificial-intelligence/agentic-ai/>

3. AI Risk Management and Security Planning

Deploying AI in engineering requires a **rigorous risk and security plan** to protect public safety, ensure ethical practices, and safeguard company interests. In this module, we identify key categories of AI-related risk, from technical failures to ethical pitfalls, and outline strategies for **risk mitigation, security controls, and safe design**. Proactively addressing these concerns is essential to developing a **trustworthy AI system** and meeting professional and regulatory obligations.

3.1 Identifying Risks of AI in Engineering

Technical Risks: AI systems can fail in unpredictable ways. For example, a structural design agent might produce an unsafe design if it extrapolates beyond its training data or if the input parameters are outside its learned range. Machine learning models are susceptible to **errors** or inaccuracies when facing scenarios not well-covered by training data. Without proper validation, an AI might overlook critical safety requirements, causing design flaws. **Robustness** is a concern, how does the AI perform under stress or when conditions change? Engineers must consider worst-case outcomes of AI decisions (e.g., what’s the consequence of a false-negative in a flaw detection system?).

Data Quality & Bias: AI is highly dependent on data, and bad data can lead to bad decisions. If an AI training dataset has gaps or errors, the model may learn incorrect patterns. Bias in data (e.g., under-representing certain scenarios or historical biases in failure data) can lead to biased or unsafe recommendations⁶³. For instance, an AI scheduling tool trained only on past projects that had ample staffing might underestimate timelines for a smaller team. *To mitigate data risks:* perform thorough data cleaning and bias checks, ensure data is current and representative of expected conditions, and plan for ongoing updates to the model as new data arrives^{64 65}.

Cybersecurity Threats: AI tools introduce new attack surfaces. An adversary could attempt to feed malicious data to an AI (called **data poisoning**) to skew its outputs. They might exploit vulnerabilities in an AI agent that has access to company systems. Additionally, generative AI agents that connect to external tools (APIs, databases) need to be safeguarded against unauthorized access. Treat your AI system as a critical piece of IT infrastructure: secure its integrations, authenticate data sources, and prevent unauthorized use. Also consider **adversarial examples**, specially crafted inputs that fool AI models (common in image recognition systems), and design the system to detect or be robust to them as feasible.

⁶³<https://www.multimodal.dev/post/ai-kpis>

⁶⁴<https://www.multimodal.dev/post/ai-kpis>

⁶⁵<https://www.multimodal.dev/post/ai-kpis>

Ethical and Legal Risks: AI decisions can raise **ethical issues**. For example, an AI might inadvertently violate privacy by exposing sensitive data, or create unfair outcomes (like disproportionately favoring certain design choices due to bias in objectives or data). There may also be **liability concerns**: Who is responsible if the AI causes an error in a design that leads to a failure? Engineers must ensure AI outputs comply with all relevant regulations, standards, and ethical guidelines. Professional engineering ethics require **holding public safety paramount**; thus any AI's actions or recommendations must be subject to that same standard. Moreover, emerging AI-specific regulations (e.g., requirements for transparency or human oversight) must be monitored and complied with as they develop.

Reputational Risk: High-profile failures of AI can damage an engineering firm's reputation. A well-known example is when an AI recommendation system or algorithm is found to be biased or makes a mistake with costly consequences, clients and the public may lose trust in the company's capabilities. Therefore, part of risk management is planning how to **communicate about the AI's role** and having contingency plans to quickly correct any mistakes. Transparent communication can mitigate fear and build trust.

3.2 Risk Mitigation Strategies and Safe Design

Building on the identified risks, this section outlines how to create a **risk and security plan** before deploying AI:



Figure 4: AI Risk Management

- Follow a Formal Risk Management Framework:** Utilize frameworks like the **NIST AI Risk Management Framework (RMF)** to systematically *map, measure, manage,* and *govern* AI risks. The NIST AI RMF (version 1.0 released 2023) provides guidelines for identifying risks across **accuracy, reliability, security, explainability, privacy, and bias**, and for implementing controls and ongoing monitoring^{66 67}. Embracing such frameworks helps ensure no major risk category is overlooked and aligns your program with industry best practices.
- Incorporate Safety from the Start:** As recommended in industry guides, safety should be **built into the AI design** (not added as an afterthought). For agentic AI, **implement multiple safety layers**⁶⁸:
 - Technical Guardrails:* hard-code constraints to prevent the AI from taking certain actions outside its authority. For instance, an agent should **fail-safe** (stop and alert) if it encounters conditions outside its training or if it is about to exceed defined safety limits (like suggesting a design violating code limits).

⁶⁶<https://www.oracle.com/artificial-intelligence/agentic-ai/>

⁶⁷<https://www.oracle.com/artificial-intelligence/agentic-ai/>

⁶⁸<https://www.oracle.com/artificial-intelligence/agentic-ai/>

- *Decision Thresholds for Human Oversight:* Define conditions under which the AI must defer to a human. For example, an AI can be allowed to auto-design a standard part, but if an unusual scenario arises or a critical decision is needed, it should request human review and approval⁶⁹.
- *Monitoring & Logging:* Have the system log its actions and decisions. This **traceability** allows engineers to audit what the AI did and why, which is crucial for transparency and debugging issues. It also supports compliance, as some regulations might require records of AI decision-making.
- *Gradual Autonomy Increase:* Initially limit the AI’s autonomy and give it narrow tasks. As confidence in its performance grows through testing and pilot use, you can widen its scope or let it operate with less supervision. This approach, analogous to supervising a new engineer hire, ensures the AI “earns trust” over time⁷⁰. In early deployments, choose **low-risk tasks** for the agent so any mistakes are low-consequence and easily corrected⁷¹.
- **Security Measures:** Work with your IT/cybersecurity team to **harden the AI system**:
 - Ensure the AI and its data pipelines have proper access controls and encryption.
 - If using external AI services or APIs, vet their security and **data privacy policies** (e.g., ensure no sensitive data is shared without safeguards).
 - Protect training data and models, as they can be targets for theft or tampering.
 - Plan for regular security audits of AI components just as you would for other critical software.
 - If the AI agent can execute actions (like controlling a system or making purchases), implement strict authentication and approval steps to avoid malicious misuse.

⁶⁹<https://www.oracle.com/artificial-intelligence/agentic-ai/>

⁷⁰<https://www.oracle.com/artificial-intelligence/agentic-ai/>

⁷¹<https://www.oracle.com/artificial-intelligence/agentic-ai/>

- **Bias and Ethics Checks:** Integrate procedures to catch and correct biases or unethical behavior:
 - Use **diverse test scenarios** to probe whether the AI's recommendations unduly favor or disfavor certain outcomes.
 - Include stakeholders during testing to identify any outputs that seem questionable or misaligned with company values or social responsibility.
 - Establish an ethical review as part of the development lifecycle, similar to a design review or safety review.
- **Validation and Testing:** Before full deployment, validate the AI extensively:
 - Use historical cases (if available) to see if the AI would have made correct decisions or designs.
 - Run simulations or sandbox tests: e.g., test a design agent on past project data to ensure it meets all requirements and produces expected results.
 - Consider **stress-testing** the agent by giving extreme or adversarial inputs to ensure it handles them safely (for example, test what a scheduling AI does if all inputs are wildly outside normal ranges, it should respond reasonably or flag issues, not produce nonsense).
 - If possible, perform a trial on a small, controlled project where experienced engineers can double-check every AI decision.
- **Safe Failure and Recovery:** Plan how the system will fail safely. In engineering, it's understood that no complex system is risk-free, what matters is how failures are managed. Ensure that if the AI malfunctions or produces an alert, the system **defaults to a safe state** (e.g., requires human confirmation, or triggers a predefined safe response). Develop an incident response plan for AI errors: how will you detect them, who will respond, how to communicate with clients or authorities if needed.

By systematically addressing these risks, an engineering firm can **confidently deploy AI agents** knowing that they have appropriate safeguards in place. The goal is to **reap AI's benefits (speed, efficiency, insight) while minimizing the chances of failures or unintended consequences.**

4. Implementing AI in Engineering Design Processes

With a solid plan and risk controls defined, the next step is to **implement AI into your engineering workflows**. This module provides best practices on integrating the AI agent into design and operational processes. Key focus areas include running pilot projects, training your team, iterating on the AI's integration, and learning from real-world case studies of AI in engineering design.

4.1 Pilot Projects and Gradual Integration

Rather than a sudden, large-scale deployment, begin with **pilot projects**. A pilot is essentially a trial run of the AI agent in a controlled setting, for example, using the AI on one project or a subset of tasks, to evaluate its performance and work out kinks. Pilots allow you to **start small, learn, and adjust** before scaling up⁷².

Key steps for an AI pilot in engineering:

- *Set Clear Goals:* Define what the pilot should accomplish. For instance, “Use the AI scheduling assistant on a single project to reduce scheduling conflicts and measure its impact on on-time task completion.” Use **SMART goals** (Specific, Measurable, Achievable, Relevant, Time-bound) so you can clearly tell if the pilot succeeded⁷³.
- *Limit Scope:* Keep the pilot focused. Only integrate the AI into one project phase or one team initially. A narrow scope makes it easier to manage and evaluate results, and limits potential disruptions.
- *Baseline and Monitor:* Before starting, record baseline metrics (e.g., current time spent on the task, current error rate or performance). Then let the AI assist, and measure the **same metrics** again. This will show the effect of the AI, for example, a reduction in design iterations, or time saved in generating reports.
- *Gather Feedback:* Collect qualitative feedback from the engineers and staff who interact with the AI. Did it actually help their work? Was it user-friendly? Did it produce any confusing or incorrect outputs that need addressing? Early user feedback is invaluable for improving both the AI's functionality and its acceptance.
- *Document Results:* Keep a detailed record of what happened in the pilot: the context, data used, outcomes achieved, issues encountered, and how they were resolved. This information will support refining the AI and also help in creating organizational knowledge for future projects⁷⁴.

After a successful pilot, you can plan to **scale up** gradually: integrate the AI into more projects or additional departments, expand its functionality, or allow greater autonomy as confidence grows.

⁷²https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁷³https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁷⁴https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

4.2 Integrating AI into the Engineering Workflow

Integrating AI requires carefully melding the agent into existing processes. Consider the typical **engineering design workflow** (which may include requirements definition, conceptual design, detailed design, analysis/simulation, review, testing, implementation, etc.) and identify where the AI will fit in:

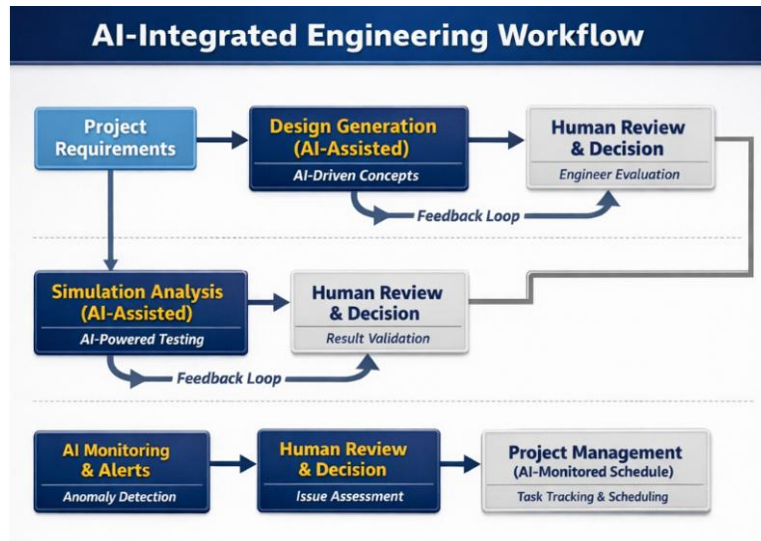


Figure 5: AI Integration in Engineering Workflow

- **During Conceptual Design:**

AI tools (especially generative design or simulation agents) can be used to generate preliminary design options and perform quick feasibility analyses. For example, an architectural design agent might produce dozens of building layouts that satisfy zoning, structural constraints, and energy efficiency criteria, giving human architects a rich set of options to consider⁷⁵. Engineers should learn how to incorporate these AI-generated options into their normal conceptual design reviews. *New step:* include an “AI-generated alternatives review meeting” early in design phases.

- **During Detailed Design and Analysis:** AI agents can automate complex analyses. In civil or mechanical engineering, an AI might automatically optimize certain parameters (like the shape of a component for stress distribution) using iterative algorithms or ML models. In electrical engineering, an AI might help with component placement and routing on a circuit board, balancing performance and cost. The **implementation challenge** is ensuring the AI’s output can be fed into your existing tools, e.g., if the AI designs a component shape, can it be imported into your CAD software for refinement? A solution is to use AI that integrates via APIs or plugins with popular engineering software, or to write translators that convert AI outputs into the needed format⁷⁶.

⁷⁵https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁷⁶https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

- **Project Management and Scheduling:** Integrating an AI scheduling assistant or risk analysis agent means adjusting how project managers work. For example, if the AI predicts risks in a project schedule, a process should be in place for the team to review those predictions in regular meetings. Perhaps the AI is configured to send weekly risk reports to project leads. **Standard Operating Procedures (SOPs)** should be updated to reflect these new steps⁷⁷.
- **Quality Control and Testing:** If an AI agent helps generate designs or analyses, incorporate additional QA steps to verify AI outputs. For instance, if an AI writes a section of an engineering report or code, have a human review it (at least initially) as part of the QA process. Organizations may define **checklists for reviewing AI-generated content** to ensure it meets quality and safety standards.
- **Feedback Loop:** Integrate a mechanism for continuous improvement. After each project that uses the AI, hold a brief retrospective focusing on the AI’s contribution: Did it meet expectations? Were there any near-misses or mistakes? Use this to refine both the AI (perhaps retraining it on new data) and the processes around it. *Treat the AI as a member of the team*; it needs performance reviews and training updates too.

Change Management: Introducing AI may change job roles and workflows. Manage this change proactively:

- *Training for Staff:* Provide training sessions for engineers and project managers on how to use the new AI tools effectively. This includes not just which buttons to click, but understanding the tool’s limitations and how to interpret its outputs. Emphasize that understanding the context is still the engineer’s responsibility; AI is there to assist with number crunching and routine tasks, but human judgment remains vital^{78 79}.
- *Process Documentation:* Update process documentation, design manuals, or project execution plans to include the AI-related steps. Clarity here helps institutionalize the new workflow so it’s repeatable and transparent.
- *Champion and Support:* Ensure the AI project team or an “AI champion” is available to support others, answer questions, and troubleshoot issues as teams start using the AI in their day-to-day work^{80 81}. Consider periodic check-ins or an internal forum for users to share tips or concerns about the AI tool.

⁷⁷https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁷⁸https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁷⁹https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁸⁰https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁸¹https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

- *Cultural Acceptance*: Celebrate successes where the AI contributed (e.g., highlight a project where AI helped save time or caught an error) to build momentum and buy-in. At the same time, be honest about any issues and how they were resolved, this keeps trust. The goal is an “**AI-ready**” culture that sees AI as a beneficial innovation. Leadership should reinforce that continuous learning (including learning from AI) is valued in the company⁸².

⁸²https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

4.3 Case Study Examples

To illustrate effective AI integration, consider these brief case studies across different engineering domains:

- **Case Study 1: Structural Engineering Design Optimization**, *A civil engineering firm* implemented a generative design AI tool to assist in optimizing bridge designs. The AI was given the goal to minimize material cost while meeting all safety and code requirements. During a pilot project, the agent generated dozens of design variants for a pedestrian bridge. Engineers reviewed the top suggestions and found a novel truss configuration that reduced steel usage by 15% while meeting load requirements. The AI's ability to autonomously search the design space yielded creative solutions that the team might not have considered, demonstrating the value of agentic AI in conceptual design. Engineers carefully verified the AI's suggestions through conventional analysis to ensure safety, and ultimately implemented the best design, resulting in cost savings for the client.
- **Case Study 2: Predictive Maintenance in Mechanical Systems**, *A manufacturing company* deployed an AI agent to monitor machine performance and predict maintenance needs. The agent autonomously collected sensor data (vibration, temperature, RPM) from dozens of industrial machines (perception), analyzed it using a predictive ML model (reasoning), and took actions such as scheduling maintenance or alerting technicians when certain thresholds were exceeded (action). Over six months, the agent correctly predicted 80% of equipment failures in advance, reducing unplanned downtime by 30%. By integrating this agent into their maintenance workflow, the company moved from a reactive maintenance approach to a proactive one. They mitigated risk by having the agent only schedule maintenance (a low-risk action) and always notify a human supervisor for approval, thus combining AI efficiency with human judgment for critical decisions.
- **Case Study 3: AI in Construction Management (AEC Industry)**, *An engineering/construction firm* introduced an AI scheduling assistant on a large construction project. The agent analyzed the project plan and continuously compared planned vs. actual progress using data from on-site reports and IoT sensors. It proactively identified potential delays, for example, noticing that a critical shipment of materials was trending late based on supplier data, and suggested resequencing some tasks to avoid idle time. The agent also monitored safety compliance via drone footage (using computer vision to check for PPE usage and hazards). During the pilot, project managers received daily recommendations from the AI. Not all suggestions were used, but some proved valuable (such as reallocating crews while waiting for materials, which saved days of delay). The firm formalized this process by having the AI's report reviewed in the morning coordination meeting each day. This case highlighted the importance of integrating AI recommendations into existing decision-making structures, and after success on one project, the firm expanded the use of the AI assistant to other projects.

By examining these examples, engineers can see how agentic AI might look in practice, working *in concert with human expertise*. Common themes include starting with a focused application, ensuring results are verified, and gradually expanding AI's role as confidence increases. In all cases, the results of AI integration were *measured* (e.g. reduction in cost or downtime) to demonstrate clear value.

5. Developing KPIs for AI Performance and Impact

To ensure an AI initiative is successful and accountable, engineering firms must establish **Key Performance Indicators (KPIs)** and metrics for their AI systems. This module explains how to define and use KPIs to **measure the AI’s performance, its effect on engineering outcomes, and ongoing alignment with business goals**. Well-chosen KPIs enable data-driven evaluation of your AI project’s success and support continuous improvement.



Figure 6: Key Performance Indicators (KPIs) for Evaluating AI Performance and Impact

5.1 Why KPIs Matter for AI Projects

In traditional engineering projects, metrics like schedule adherence, budget variance, number of design iterations, or safety incident counts are tracked to gauge project performance. Similarly, when you introduce an AI agent into the process, you need to track its contributions and behavior. **“You can’t manage what you don’t measure,”** as the saying goes⁸³. KPIs for AI serve several purposes:

- Evaluate Effectiveness:** Determine if the AI is doing what it’s supposed to do. Is the design optimization agent actually reducing design time or improving quality? Is the predictive maintenance agent accurately predicting failures? KPIs help answer these questions objectively.
- Align with Business Goals:** By tying AI performance to business outcomes (like cost savings, efficiency gains, safety improvements), you ensure the AI project stays focused on meaningful results rather than just technical performance.
- Guide Adjustments:** Continuous measurement allows the team to make data-driven adjustments. If an AI isn’t meeting a certain KPI (say, its accuracy is below target), you can decide to retrain the model, provide more data, or refine its algorithms. If the AI is meeting its technical KPIs but not improving the business metric (e.g., no cost savings despite high model accuracy), that might mean the chosen application isn’t as impactful as expected and the strategy needs revisiting⁸⁴.
- Demonstrate Value and Accountability:** For stakeholders and regulators, KPIs are evidence that the AI is under control and delivering value. This is important for justifying further investment and for meeting any compliance obligations (for instance, if audited, you can show that you monitor the AI’s error rates, safety compliance, etc.).

⁸³<https://cloud.google.com/transform/gen-ai-kpis-measuring-ai-success-deep-dive>

⁸⁴<https://cloud.google.com/transform/gen-ai-kpis-measuring-ai-success-deep-dive>

5.2 Categories of KPIs for AI

When developing KPIs for an AI in engineering, consider a mix of **technical performance metrics** and **business/operational metrics**:

- **Accuracy and Quality Metrics:** If the AI makes predictions or classifications, use metrics like **accuracy**, **precision/recall**, or **error rate** to measure its output quality⁸⁵ ⁸⁶. For example, an AI flaw detection system might track the percentage of defects correctly identified (true positives) and false alarms raised (false positives). For a design-generating AI, quality might be measured by how often its suggestions pass human review or meet all requirements.
- **Efficiency and Throughput:** Measure the **speed** or throughput of the AI's tasks. How many designs per hour can it generate? How quickly does it analyze a data set and produce results? Also track **uptime** and reliability, especially if the AI is part of a live system, you might use metrics like system availability (percentage of time the AI system is operational) and response time or latency for real-time agents⁸⁷. If the AI is too slow to be useful or frequently unavailable, those issues need addressing.
- **Adoption and Usage:** For internal AI tools, gauge how widely and effectively they are being used by the engineering team. Metrics might include the **number of projects or teams using the AI**, the frequency of use, or percentage of project tasks that involve the AI. Low adoption could indicate usability issues or lack of training, whereas increasing adoption (paired with positive feedback) shows growing trust in the tool.
- **Outcome/Impact Metrics:** These link the AI to business or project outcomes. Examples:
 - **Time Savings:** Reduction in time taken for a task due to AI (e.g., “The AI scheduling tool cut the weekly schedule update process from 5 hours to 2 hours”).
 - **Cost Savings or ROI:** Financial impact of AI, e.g., money saved through efficiency or error reduction. This could be measured per project or annually. For instance, if an AI prevented a costly design error, you might quantify the avoided rework cost.
 - **Quality and Safety Improvements:** Metrics like number of defects found (by AI vs. previously), reduction in safety incidents or non-compliance events after implementing an AI monitoring system.
 - **Client Satisfaction or Deliverable Quality:** If applicable, measure client feedback or quality scores on projects with AI assistance vs without.

⁸⁵<https://www.multimodal.dev/post/ai-kpis>

⁸⁶<https://www.multimodal.dev/post/ai-kpis>

⁸⁷<https://cloud.google.com/transform/gen-ai-kpis-measuring-ai-success-deep-dive>

- **Compliance and Ethical Metrics:** If you have internal or external requirements, translate them into metrics. For example, **percentage of AI decisions reviewed by a human** (to ensure oversight), or **number of bias incidents detected in AI outputs** (with a goal of zero). If the organization has an AI ethics checklist, track completion of that for each deployment.

When selecting KPIs, choose only those **most relevant to the AI’s purpose and your goals**. Aim for a handful of clear metrics rather than an overwhelming list. Each KPI should have a defined target or benchmark (e.g., “90% accuracy on identifying defects” or “reduce design cycle by 25% in six months”).

Also be mindful that improving one metric might affect another. For instance, an AI-driven design might decrease material cost (positive) but could increase design complexity or verification time (negative). Monitor for any unintended side effects in your metrics and adjust accordingly^{88 89}.

5.3 Monitoring and Continuous Improvement

After defining KPIs, implement a process to **monitor them regularly**. This could involve dashboards or automated reports. Many AI platforms provide tools for tracking model performance over time. You might integrate alerts for when a metric goes out of expected bounds (for example, if the AI’s error rate spikes above a threshold, notify the team immediately).

Use the KPI data in recurring reviews, perhaps as part of project post-mortems or monthly quality meetings, to discuss the AI system’s performance. Questions to consider: Are we seeing the promised benefits? If not, why? Is it a technical issue (e.g., model needs retraining) or a deployment issue (e.g., not everyone is using the tool effectively)? Continuous monitoring is also an **early warning system for risks**: a degrading KPI might signal model drift, data pipeline problems, or user workarounds that indicate the AI output isn’t meeting their needs^{90 91}.

Finally, update your KPIs over time as needed. As your AI matures, you may add more ambitious metrics (for example, once an AI achieves reliability, you might start measuring how much *higher-level value* it adds, like impact on profit or new business won). The KPI framework itself can evolve with the organization’s AI journey.

By rigorously tracking KPIs, engineering firms can ensure their agentic AI projects stay on track and deliver tangible benefits. This data-driven approach also supports transparency and accountability, critical components when introducing transformative technologies into professional practice.

⁸⁸<https://cloud.google.com/transform/gen-ai-kpis-measuring-ai-success-deep-dive>

⁸⁹<https://cloud.google.com/transform/gen-ai-kpis-measuring-ai-success-deep-dive>

⁹⁰<https://www.oracle.com/artificial-intelligence/agentic-ai/>

⁹¹<https://www.oracle.com/artificial-intelligence/agentic-ai/>

6. Policy and Governance for AI Implementation

The final module addresses the creation of an **AI policy and governance framework** for your engineering firm. Successful AI adoption isn't just about technology, it also requires **organizational policies, ethical guidelines, and oversight structures** to manage how AI is used. A robust governance plan ensures that the deployment of agentic AI aligns with legal requirements, industry standards, and the company's values, thereby protecting the public and the organization.

6.1 Developing an AI Policy

An **AI Policy** is a formal document (or set of documents) that outlines how your company will utilize AI. It typically covers:

- **Scope of AI Use:** Define what types of AI applications are permitted and in what areas of your business. For example, a policy might state that “AI may be used to assist in design and analysis tasks, but final engineering judgments must be made or reviewed by a licensed professional engineer.” This ties into regulatory requirements that **licensed engineers take responsibility for engineering decisions**, ensuring AI doesn't inadvertently override human judgment in critical matters.
- **Data Governance and Privacy:** The policy should mandate proper handling of data used by AI. It might include guidelines such as: ensuring data is obtained and used in compliance with privacy laws and client agreements, prohibiting the input of confidential project data into external AI tools without clearance, and requiring anonymization of sensitive information where possible.
- **Intellectual Property (IP):** Clarify ownership of AI-generated content or designs. For instance, if an AI creates a novel design, does the company treat it as company IP? Address how using third-party AI (like cloud-based services) might impact IP , some cloud AI providers claim rights or usage of data fed into them, which can conflict with client confidentiality. The policy should require vetting of such terms before using any external AI service.
- **Risk Management & Quality Assurance:** State the requirement for risk assessments and validations of AI (as covered in Module 3). The policy can require that any AI tool undergoes a review and approval (by an AI governance committee or similar) before being used on live projects, and that critical AI systems have human oversight. Include reference to following established industry guidelines for “**Responsible AI**” or frameworks like NIST's AI Risk Management Framework for comprehensive coverage of safety, fairness, and reliability^{92 93}.

⁹²<https://www.oracle.com/artificial-intelligence/agentic-ai/>

⁹³<https://www.oracle.com/artificial-intelligence/agentic-ai/>

- **Ethical Principles:** Many organizations choose to articulate principles for AI usage , for example, commitments to fairness, transparency, accountability, and privacy. These principles act as a north star for all AI projects. For an engineering firm, this might also include a commitment that AI will be used in ways consistent with the **NSPE (National Society of Professional Engineers) Code of Ethics**, reinforcing that **AI will not be deployed in ways that compromise public safety or well-being**.
- **Compliance with Laws and Standards:** Even though AI-specific regulations are still evolving, your policy should acknowledge current laws (like data protection regulations, or sector-specific rules if you're in a regulated industry such as healthcare or transportation). It should commit to following any future applicable laws (for example, the EU *AI Act* if working internationally, or guidelines from U.S. agencies for AI in medicine, transportation, etc.). Also consider referencing standards from professional bodies (like IEEE's guidelines on Ethics of AI, or ISO/IEC standards on AI management) to show alignment with global best practices in the absence of strict regulations.

When drafting the AI policy, involve both technical experts (who understand the AI) and non-technical leadership (who understand business and legal implications). **Engage legal counsel** to review the policy and address liability, contract terms, and regulatory compliance⁹⁴. For example, contracts with clients might need new clauses clarifying the use of AI in deliverables, and contracts with AI vendors need clauses on data use and liability. Given the still-changing legal landscape, plan to **regularly update the policy** as new regulations and case law emerge⁹⁵.

⁹⁴https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

⁹⁵https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

6.2 Governance Structure and Roles

Implementing policy is not a one-time effort, it requires ongoing **governance**. Companies should establish formal **AI governance structures** to oversee AI initiatives. Depending on the size of the firm, this could be an **AI Steering Committee**, a designated **AI Officer**, or part of an existing risk management committee. The ACEC recommends, for larger firms, forming an AI committee with representatives from various departments (IT, engineering, operations, legal, etc.)⁹⁶. Smaller firms might simply assign an “AI Champion” or rely on existing leadership roles, but they should still have clear accountability in place.



Figure 7: AI Governance Structure: Roles and Organizational Oversight

The responsibilities of AI governance typically include:

- **Approving AI Projects:** Evaluating proposed AI use cases for alignment with strategy and risk profile. The governance body can ensure that a proper risk assessment (as in Module 3) is done before a project proceeds.
- **Monitoring AI Performance and Compliance:** Overseeing the KPIs and reports (Module 5) for deployed AI systems. The governance team should regularly review whether the AI is meeting its targets and adhering to policies (e.g., checking that no policy violations or near-misses have occurred, such as an AI almost releasing unverified output).
- **Policy Updates and Training:** Keeping the AI policy up-to-date. For example, if a new regulation or standard is released, the governance team updates the policy and communicates changes. They may also be in charge of organization-wide **AI ethics training** or awareness programs so that all staff understands the do’s and don’ts of AI use.
- **Incident Response:** In case of an AI-related incident (say the AI gives a dangerous recommendation that was caught, or worse, an error that wasn’t caught in time), the governance group must investigate what went wrong, ensure proper disclosure if required, and decide on corrective actions (like adjusting the system or retraining staff). Having a predefined incident response plan for AI is as important as it is for other aspects of engineering emergencies.

⁹⁶https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

- **Continuous Improvement:** Steering the strategic direction of AI in the company, deciding when to scale successful pilots, prioritize new AI opportunities, and allocate resources. The governance body ensures that the company’s AI adoption roadmap (see final section) moves forward responsibly and is aligned with business strategy.

Documentation and Transparency: Good governance involves documentation of AI systems (data sources, algorithms, versions, known limitations) and decisions made about them. This documentation aids in transparency. In regulated environments, you might need to demonstrate to regulators how your AI is controlled; having a paper trail via your governance group’s records is invaluable.

Finally, governance should foster a culture of **responsibility and ethics** around AI. This means encouraging employees to speak up if they see something concerning in an AI’s behavior, and protecting them (whistleblower style) if they raise ethical issues. It also means rewarding teams who integrate AI in ways that improve outcomes and uphold values, as positive reinforcement for responsible innovation⁹⁷.

In summary, **Policy and Governance** provide the necessary oversight to accompany the technical work. They ensure that as agentic AI is woven into engineering tasks, it is done in a controlled, ethical, and legally compliant manner. This safeguards not only the public and end-users of engineering projects, but also the company’s reputation and engineers’ professional licenses.

⁹⁷https://www.acec.org/wp-content/uploads/2025/10/AI-Adoption-in-AEC-Summary_Primer_October_2025.pdf

7. Developing a Safe AI Implementation Action Plan

Bringing together all the previous modules, we can outline a step-by-step **action plan for implementing agentic AI in an engineering company**. This serves as a roadmap that your team can follow to ensure a methodical and safe adoption of AI:

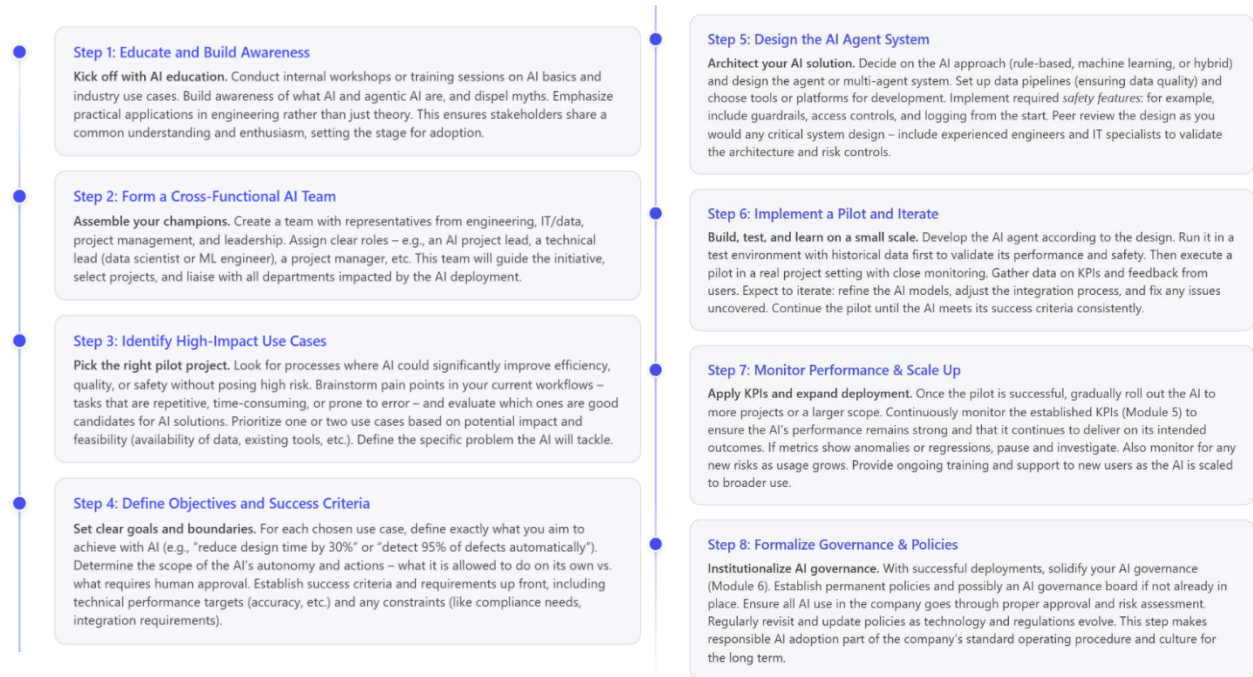


Figure 8: Safe AI Implementation Plan

Following these steps provides a **safe and structured pathway** for bringing agentic AI into your engineering practice. By educating your team, starting with the right project, diligently managing risks, and embedding governance, you create a balance between **innovation and safety**. The result is an AI-enhanced engineering capability that can improve efficiency and outcomes while maintaining trust, compliance, and ethical standards.

8. Conclusion and Next Steps

Artificial intelligence, especially agentic AI, represents the next big leap in engineering practice, allowing for greater automation of complex tasks and offering new insights that can transform how projects are designed, executed, and managed. **Agentic AI systems can act as tireless assistants or even semi-autonomous team members**, handling routine work and synthesizing information to support human engineers in making better decisions. This course has provided an overview of AI fundamentals and a deep dive into the key considerations for introducing such technology in a professional engineering context.

It is crucial to remember that **the goal is not to adopt AI for its own sake**, but to solve real engineering problems more effectively. By following the safe implementation process outlined, from conceiving valuable use cases and carefully designing agents, through managing risks, integrating into workflows, measuring performance, and instituting strong governance, an engineering firm can confidently leverage AI **while upholding the paramount values of safety, quality, and ethics**.

Engineers are, at their core, problem solvers and innovators. AI is another powerful tool in the modern engineer's toolkit. Those who learn to use it responsibly will be able to **augment their capabilities, drive efficiency, and push the boundaries of innovation** in their field, all while safeguarding the public and the integrity of the profession.

As a final step in this course, the following section provides an **assessment quiz** to test and reinforce your understanding of the material. Good luck, and thank you for engaging in this important professional development journey!

Summary:

This **three-hour online professional-development course** introduces licensed engineers to agentic AI, offering a safe framework for integrating AI into engineering practice. It covers definitions, design principles, risk mitigation, workflow integration, performance measurement and governance. Throughout, the course emphasizes professional responsibilities, safe, ethical and compliant application of AI, and draws on industry research (cited below).

Key Topics & Structure

1. Introduction to AI and Agentic AI

- **AI fundamentals:** Definitions of artificial intelligence, machine learning and generative AI. Generative design tools can explore **thousands of design permutations** based on specified goals and constraints, helping engineers speed exploration and start with optimized options. These tools can reduce product development time by **30, 50 %** and cut costs up to **20 %** while reducing weight by 10, 50 %. [[nist.gov](https://www.nist.gov)]
- **Agentic AI:** Differentiates between traditional generative or predictive AI (reactive tools) and agentic AI, which autonomously plans and performs actions towards goals. Discusses real-world applications such as predictive maintenance and project management assistants.
- **Industry adoption:** Within the architecture, engineering and construction sector, **63 %** of member firms have an AI strategy or are developing one. However, many firms still feel unprepared, highlighting the need for education and governance. [[designnews.com](https://www.designnews.com)] [[nist.gov](https://www.nist.gov)]

2. Conceiving & Designing AI Agents

- **Problem selection:** Identify high-impact, data-driven tasks or design processes where AI can improve efficiency or quality, starting with pilot use cases. Ensure alignment with business goals.
- **Goal & scope definition:** Set clear objectives for the agent, boundaries for its autonomy (advisory vs. semi-autonomous), and integration requirements with existing data and software.
- **Design process:** Choose appropriate AI methodologies, rule-based, machine learning or hybrid. Decide whether to use a single agent or multi-agent system, and design guardrails to maintain safety. Involve cross-functional teams (engineering, IT, management) early to address technical and organizational needs.

3. AI Risk and Security Planning

- **Risk categories:** Address technical risks (model errors, training-data bias), data quality concerns, cyber-attack vectors and ethical/legal issues. Bad data or unrepresentative training sets can cause biased or unsafe decisions. [\[nist.gov\]](https://www.nist.gov)
- **Risk management frameworks:** Introduce the **NIST AI Risk Management Framework**, which is a voluntary guideline to manage risks and incorporate **trustworthiness** in AI design, deployment and use. [\[report.md\]](#) | [PowerPoint](#)
- **Mitigation strategies:** Emphasize safe-fail design, human-in-the-loop approval for high-stakes actions, continuous testing, model and data security, bias assessment, and regular audits.

4. Implementing AI in Design Processes

- **Pilot projects:** Begin with limited-scope pilots; measure baseline metrics and gather user feedback. Generative design case studies illustrate how AI can quickly produce many feasible design alternatives, saving time and discovering innovative forms. [\[nist.gov\]](https://www.nist.gov)
- **Workflow integration:** Explain how to embed AI steps into conceptual design, detailed analysis, project scheduling, and quality-control processes. Discuss training staff, updating standard operating procedures, and ensuring human review of AI outputs.
- **Case studies:** Summarize examples such as structural design optimization, predictive maintenance and project management assistants to illustrate real-world benefits and challenges.

5. Developing KPIs for AI Performance

- **Purpose:** Measure AI effectiveness and align performance with business and safety goals. Use metrics to guide refinement and justify investment.
- **Categories:** Technical KPIs (accuracy, precision, false-positive rate); efficiency metrics (time savings, throughput); adoption metrics (usage rates); impact metrics (cost savings, quality improvement); and compliance/ethical metrics (percentage of AI-generated decisions reviewed by engineers).
- **Monitoring:** Use dashboards to track metrics over time and detect model drift or implementation issues. Combine quantitative data with qualitative feedback.

6. Policy & Governance Plans

- **AI policy:** Develop organizational policies covering allowable uses of AI, data governance and privacy, intellectual-property rights, risk assessment procedures and adherence to professional ethics.
- **Governance structure:** Establish an AI steering committee or assign AI responsibility within existing leadership. The committee should evaluate requests, monitor KPIs, ensure compliance with the AI policy, and update policies as regulations evolve.
- **Regulatory alignment:** Acknowledge emerging regulations (e.g., EU AI Act) and professional standards; ensure licensed engineers maintain accountability for AI-assisted decisions.

7. Implementation Roadmap

A step-by-step action plan synthesizes earlier modules:

1. **Educate** staff about AI fundamentals and industry use cases.
2. **Form a cross-functional AI team** and secure executive sponsorship.
3. **Identify pilot use cases** where AI can deliver measurable benefits.
4. **Define goals, scope and success criteria** for each use case.
5. **Design and develop the AI agent** with safety guardrails and risk controls.
6. **Implement pilots**, collect data and iterate.
7. **Scale up** successful AI solutions and monitor KPIs.
8. **Formalize governance** and update policies as AI adoption grows.

Conclusion

Adopting agentic AI can enhance the productivity, creativity and safety of engineering practice when implemented responsibly. By following a structured process, identifying appropriate opportunities, designing with safety in mind, managing risks through frameworks like NIST's AI RMF, integrating AI into workflows, measuring performance and establishing governance, engineers can unlock AI's benefits while upholding professional obligations and protecting the public. This course equips engineers with the knowledge and frameworks needed to safely implement agentic AI in their organizations.