Artificial Intelligence Glossary for AEC Leadership

Executive Summary

Artificial intelligence is introducing a new operating language into architecture, engineering, and construction. Terms such as *agentic AI*, *retrieval-augmented generation*, *digital twin*, and *model governance* are moving from technology discussions into decisions about project delivery, investment, risk, workforce strategy, contracts, and client services.

The purpose of this glossary is not to turn executives into data scientists. It is to give leaders a shared vocabulary for asking better questions, evaluating proposals, approving investments, managing risk, and communicating expectations across business, technology, legal, operational, and project teams.

A common vocabulary is itself a governance mechanism. The National Institute of Standards and Technology recommends that organizations define key AI terms, clarify human roles, connect AI governance to existing organizational controls, and establish practices for mapping, measuring, and managing AI risk. ISO/IEC 42001[1] Similarly, it provides a structured management-system approach to governing AI risks and opportunities.

For AEC organizations, this shared language is particularly important because AI may influence:

- Design development and option analysis
- Cost estimating and quantity takeoff
- Scheduling and resource allocation
- Safety monitoring and quality control
- Contract and specification review
- Requests for information, submittals, and document control
- Building information modeling and digital twins
- Field progress measurement
- Recruitment, performance management, and workforce planning
- Etc.

The glossary should be treated as a practical decision tool rather than a static dictionary. Each term includes an executive-level definition, an AEC context, and an explanation of why the term matters.

---

[1] ISO/IEC 42001:2023 is the world's first international standard for Artificial Intelligence Management Systems (AIMS), published by the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC). It provides a comprehensive framework for organizations to establish, implement, maintain, and continually improve the responsible development, provision, and use of AI systems.

Recommendations for Executive Use

1. Adopt the Glossary as a Firmwide Standard

Approve the glossary through an AI steering committee or equivalent governance body. Use the same definitions in **strategy documents, business cases, procurement materials, policies, training, contracts, and board communications.**

Prevent teams from using the same term to describe fundamentally different capabilities. In particular, require proposals to distinguish among copilots, agents, agentic systems, conventional automation, predictive models, and generative AI.

2. Build Use Cases

Create an AI use-case register. For each proposed or operating use case, document:

- Business problem and intended outcome
- Executive sponsor and accountable owner
- Users and affected stakeholders
- Data sources and system integrations
- Type of AI involved
- Degree of autonomy
- Required human review
- Risk classification
- Performance measures
- Vendor and contract status
- Deployment and retirement criteria

The glossary becomes more valuable when executives can connect its language to real systems and decisions.

3. Build the Terms into Use Case Proposals

Require every AI business case to answer five questions:

1. **Value:** What measurable business or client outcome will improve?
2. **Information:** What data will the system access, and is that data reliable and authorized?
3. **Reliability:** How will the organization evaluate accuracy, consistency, and limitations?
4. **Control:** What actions can the system take, and where is human approval required?
5. **Economics:** What are the full costs of integration, licensing, security, training, monitoring, change management, and vendor dependence?

This moves the conversation from “Do we need AI?” to “Which capability creates defensible value under acceptable conditions?”

4. Establish Risk Tiers

Classify applications according to consequence rather than novelty.

Lower-risk examples: Internal brainstorming, formatting, translation, meeting summaries, and first-draft administrative content.

Moderate-risk examples: Project forecasts, document classification, proposal analysis, productivity recommendations, and client-facing drafts.

Higher-risk examples: Safety decisions, engineering calculations, code compliance, contract interpretation, payment approval, personnel decisions, autonomous equipment, and outputs carrying professional liability.

Controls should increase with the consequence of an incorrect output or unauthorized action.

5. Use the Glossary as a Procurement Scorecard

Require vendors to explain their products using your firm's terminology. Procurement teams should ask:

- Is the product a copilot, an agent, or a fully automated workflow?
- Which foundation models does it use?
- Is customer data used to train any model?
- Where is data stored and processed?
- How is access controlled?
- How are responses grounded?
- Can the system provide citations or source records?
- How is accuracy evaluated?
- What guardrails restrict actions?
- How are model changes communicated?
- Can data and configurations be exported?
- What happens when the underlying model or vendor changes?
- How are incidents reported and investigated?

Clear terminology makes vague sales claims harder to hide behind a cloud of technological glitter.

6. Add AI Reviews to Existing Project and Delivery Gates

AI should not operate as a separate governance island. Incorporate appropriate reviews into:

- Technology architecture
- Cybersecurity
- Data governance
- Legal and contractual review
- Quality management
- Project startup
- Design review
- Safety management
- Procurement
- Change management
- Records retention
- Internal audit

Using existing governance structures generally produces faster adoption and clearer accountability than creating a parallel approval universe.

7. Pair Every Pilot with Baselines and Evaluations

Before beginning a pilot, record the current cost, cycle time, error rate, rework level, or service outcome. Define the minimum improvement required for continuation.

Evaluate more than speed. Useful measures may include:

- Accuracy
- Completeness
- User adoption
- Time saved
- Rework avoided
- Risk detected
- Response consistency
- User confidence
- Client experience
- Cost per transaction
- Number and severity of human corrections

A pilot should end with a decision to scale, revise, contain, or discontinue the application.

8. Make AI Literacy Role-Based

Executives need to understand investment, risk, accountability, and organizational implications. **Project professionals** need to understand approved uses, verification requirements, and information handling. **Technology teams** need deeper knowledge of architecture, security, evaluation, and monitoring.

Training should use actual AEC scenarios rather than generic demonstrations.

9. Review the Glossary Regularly

Assign an owner to maintain the glossary and record:

- New terms
- Changed definitions
- Retired terminology
- Relevant standards
- Regulatory developments
- New security threats
- Lessons from internal pilots and incidents

A controlled glossary can become part of the organization's broader AI management system.

The Top 50 AI Terms for AEC Executives

I. Foundational AI Concepts

1. Artificial Intelligence

Definition: Artificial intelligence, or AI, is the broad category of computer systems designed to perform tasks associated with human capabilities, including perception, language understanding, pattern recognition, prediction, reasoning, and decision support.

AEC context: AI can analyze drawings, classify project documents, forecast schedule delays, identify safety conditions, generate design options, or answer questions about project information.

Why it matters: Executives should distinguish AI from conventional software. Traditional software follows explicitly programmed rules; AI often generates results from patterns learned from data and therefore introduces different forms of uncertainty, risk, and oversight.

2. Machine Learning

Definition: Machine learning is a branch of AI in which systems learn patterns from historical data and use those patterns to make predictions, classifications, or recommendations.

AEC context: A machine-learning system might predict whether a project will exceed its budget based on scope, procurement status, labor availability, change-order history, and production trends.

Why it matters: The value and reliability of machine learning depend heavily on the relevance, quality, and representativeness of the data used to develop it.

3. Deep Learning

Definition: Deep learning is a form of machine learning that uses multilayered neural networks to identify complex patterns in large quantities of data.

AEC context: Deep learning supports applications such as image recognition, speech processing, automated drawing interpretation, site-camera analysis, and equipment-failure prediction.

Why it matters: Deep-learning systems can be powerful but are often computationally intensive and difficult to explain. Executives should understand when their complexity is justified by measurable business value.

4. Generative AI

Definition: Generative AI creates new content, including text, images, video, audio, code, designs, summaries, and structured data, based on patterns learned during training.

AEC context: Generative AI can draft proposals, summarize specifications, develop meeting minutes, create conceptual imagery, generate design alternatives, or prepare initial responses to RFIs.

Why it matters: Generated content may look authoritative even when it is incomplete or incorrect. Its use requires verification, information controls, and clear accountability.

5. Foundation Model <add Frontier, Open and Closed Models>

Definition: A foundation model is a large, broadly trained AI model that can be adapted to perform many different tasks.

AEC context: A general foundation model may be connected to project documents, estimating databases, design standards, or asset information to support specialized AEC workflows.

Why it matters: Many AI products are interfaces built on the same underlying foundation models. Executives should evaluate the complete solution, not simply the model or brand name.

6. Large Language Model

Definition: A large language model, or LLM, is a foundation model trained to understand and generate language by predicting patterns in text and other structured information.

AEC context: LLMs can search specifications, summarize contracts, compare proposals, draft correspondence, extract obligations, and provide conversational access to project information.

Why it matters: An LLM does not understand a project in the same way a qualified professional does. It produces statistically plausible responses and may require grounding, evaluation, and human approval.

7. Multimodal AI

Definition: Multimodal AI can process or generate more than one type of information, such as text, images, drawings, video, audio, point clouds, and numerical data.

AEC context: A multimodal system might review a site photograph, compare it with a BIM model, read an associated inspection report, and identify a potential discrepancy.

Why it matters: AEC information rarely exists in one format. Multimodal capability is therefore central to many high-value industry applications.

8. Natural Language Processing

Definition: Natural language processing, or NLP, is the use of computing methods to interpret, classify, search, translate, or generate human language.

AEC context: NLP can extract clauses from contracts, categorize RFIs, identify recurring issues in daily reports, and search project records using ordinary questions.

Why it matters: A large portion of project risk is buried in unstructured text. NLP can make that information more accessible, but it does not eliminate the need for professional interpretation.

9. Computer Vision

Definition: Computer vision enables systems to interpret information contained in photographs, video, scanned documents, drawings, and other visual media.

AEC context: Computer vision can support progress tracking, personal protective equipment detection, defect identification, material recognition, quantity measurement, and site documentation.

Why it matters: Visual AI can expand field visibility, but accuracy can be affected by lighting, camera position, obstructions, weather, and project-specific conditions.

10. Predictive Analytics

Definition: Predictive analytics uses historical and current information to estimate the likelihood of future outcomes.

AEC context: Applications include forecasting cost overruns, schedule slippage, safety incidents, equipment failure, labor demand, claims exposure, and asset-maintenance needs.

Why it matters: Predictions should inform decisions rather than be treated as certainties. Executives should ask how accurate the forecast is, what assumptions it makes, and what action follows from it.

11. Optimization

Definition: Optimization uses mathematical or computational methods to identify a preferred solution among multiple alternatives and constraints.

AEC context: Optimization can support structural design, site layout, energy performance, logistics, crew allocation, equipment utilization, procurement, and scheduling.

Why it matters: An optimized answer is only as useful as the objective it was given. A system optimized solely for cost may unintentionally compromise resilience, constructability, safety, quality, or carbon performance.

II. AI Assistants, Agents, and Workflows

12. AI Copilot

Definition: An AI copilot assists a person by recommending, drafting, summarizing, searching, or analyzing information while leaving final control with the user.

AEC context: A project-management copilot might draft meeting minutes, locate contract requirements, prepare a status summary, or suggest a response to a project issue.

Why it matters: Copilots can improve productivity without fully automating decisions. The organization must still define what information they can access and what outputs require review.

< add distinction of MCP, Harness, Artifacts, etc.>

13. AI Agent

Definition: An AI agent is a system that can interpret an objective, select actions, use software tools or data sources, and complete multiple steps toward a result.

AEC context: An agent might review incoming submittals, locate applicable specification sections, identify missing information, prepare routing recommendations, and update a workflow platform.

Why it matters: An agent can do more than generate text. Because it may take actions, access systems, or alter records, its permissions and approval boundaries must be tightly controlled.

14. Agentic AI

Definition: Agentic AI describes systems that operate with a degree of autonomy by planning, taking actions, evaluating intermediate results, and adjusting their approach.

AEC context: An agentic system could monitor project controls, investigate emerging variances, request information from connected systems, and prepare recommended interventions.

Why it matters: Increasing autonomy increases both potential value and potential consequence. Executives should require clear limits, audit trails, escalation rules, and human approval for consequential actions.

15. Intelligent Automation

Definition: Intelligent automation combines AI with workflow tools, business rules, integrations, and sometimes robotic process automation.

AEC context: It can automate repetitive processes such as invoice classification, document routing, data entry, compliance checking, and project-report compilation.

Why it matters: The largest gains often come from redesigning an entire workflow rather than adding an AI chatbot to an inefficient process.

16. Prompt

Definition: A prompt is the instruction, question, context, or data provided to a generative AI system.

AEC context: A prompt may ask an AI system to compare two specifications, identify schedule risks, summarize a contract clause, or draft a client communication.

Why it matters: Vague prompts tend to produce vague results. High-value prompts typically state the role, objective, source material, constraints, required format, and verification standard.

17. Prompt Engineering

Definition: Prompt engineering is the practice of designing and refining instructions so an AI system produces more useful, reliable, and consistent results.

AEC context: A firm may create approved prompt templates for reviewing contracts, preparing project summaries, analyzing change requests, or evaluating design narratives.

Why it matters: Prompt templates can convert individual experimentation into a repeatable organizational process. They should be tested, governed, and version controlled like other business procedures.

18. Context Window

Definition: The context window is the amount of information an AI model can consider during a single interaction or task.

AEC context: A system reviewing a large contract set may not be able to consider every drawing, specification, addendum, and correspondence item simultaneously.

Why it matters: A larger context window does not guarantee better comprehension. Organizations need document-selection, retrieval, and validation strategies for complex project records.

19. Retrieval-Augmented Generation

Definition: Retrieval-augmented generation, or RAG, retrieves relevant information from approved sources and supplies it to a generative model before the model produces an answer.

AEC context: A RAG system can answer questions using a company's standards, project specifications, lessons learned, contracts, BIM data, or operations manuals.

Why it matters: RAG can make responses more relevant and traceable, but it does not guarantee correctness. Retrieval quality, source permissions, document currency, and citation accuracy remain critical.

20. Fine-Tuning

Definition: Fine-tuning further trains an existing model using a specialized dataset so that it performs a particular task, follows a preferred style, or recognizes domain-specific patterns.

AEC context: A model might be fine-tuned to classify construction documents, interpret discipline-specific terminology, or categorize historical project risks.

Why it matters: Fine-tuning can improve specialized performance, but it introduces costs associated with data preparation, testing, security, monitoring, and model maintenance.

21. Grounding

Definition: Grounding connects an AI-generated response to specific, authoritative information rather than relying only on the model's general training.

AEC context: A grounded contract assistant should base its response on the executed agreement and cited project documents rather than on general construction knowledge.

Why it matters: Grounding improves reliability and auditability. Executives should ask whether outputs are linked to approved sources and whether users can inspect those sources.

22. Embedding

Definition: An embedding is a numerical representation of information that captures relationships in meaning or similarity.

AEC context: Embeddings can help a system recognize that "delay damages," "liquidated damages," and related contract language may concern similar concepts even when the wording differs.

Why it matters: Embeddings enable semantic search and RAG, but poorly managed embeddings may expose confidential information or return misleading matches.

23. Vector Database

Definition: A vector database stores and retrieves embeddings so information can be found according to conceptual similarity rather than exact keywords.

AEC context: It may enable conversational searches across proposals, RFIs, submittals, standards, inspection reports, and lessons-learned libraries.

Why it matters: Vector databases can become important repositories of organizational knowledge. They therefore require access controls, retention policies, source tracking, and security monitoring.

24. Application Programming Interface

Definition: An application programming interface, or API, is a controlled method by which one software system exchanges information or requests actions from another.

AEC context: APIs may connect an AI system with project-management platforms, BIM repositories, ERP systems, scheduling software, document-management systems, or field applications.

Why it matters: The business value of AI often depends on integration. Executives should understand the cost, security, reliability, and vendor-dependency implications of these connections.

25. AI Literacy

Definition: AI literacy is the ability to understand AI's basic capabilities, limitations, risks, and appropriate uses.

AEC context: Different levels of literacy are required for board members, executives, project teams, technology staff, legal counsel, and employees operating AI-enabled workflows.

Why it matters: AI literacy reduces both reckless adoption and unnecessary resistance. It helps people recognize when an output is useful, when it must be challenged, and when expert judgment is required.

III. AEC Information and Technology Environment

26. Building Information Modeling

Definition: Building information modeling, or BIM, is a process for creating, coordinating, and managing structured digital information about a built asset.

AEC context: AI may analyze BIM data for design coordination, quantity extraction, constructability, code review, progress comparison, operations, and asset management.

Why it matters: BIM provides structured context that can make AI more useful. Inconsistent modeling practices, incomplete attributes, and unclear ownership can substantially limit that value.

27. Digital Twin

Definition: A digital twin is a digital representation of a physical asset or system that exchanges information with the physical environment over its lifecycle.

AEC context: A digital twin may combine BIM, sensor data, equipment information, maintenance history, occupancy, energy use, and operational performance.

Why it matters: Digital twins can support monitoring, simulation, predictive maintenance, and asset optimization. They require a long-term information strategy, not simply a three-dimensional model.

28. Digital Thread

Definition: A digital thread is the connected flow of asset and project information across planning, design, construction, commissioning, operations, renovation, and retirement.

AEC context: It links decisions, requirements, models, documents, materials, equipment, inspections, and performance records across project phases.

Why it matters: AI is more effective when information remains connected and traceable. Fragmented records create blind spots, duplicated effort, and unreliable outputs.

29. Common Data Environment

Definition: A common data environment, or CDE, is an agreed system or process for collecting, managing, reviewing, sharing, and controlling project information.

AEC context: It may contain drawings, models, specifications, RFIs, submittals, meeting records, field documentation, and approval histories.

Why it matters: A CDE can provide authoritative sources for AI applications. Poor permissions, inconsistent naming, outdated documents, or uncontrolled duplicates can undermine those applications.

30. openBIM and Interoperability

Definition: Interoperability is the ability of different software systems to exchange and use information effectively. openBIM applies open standards and workflows to support this exchange across the built-asset lifecycle.

AEC context: It enables information created in one design, construction, geospatial, or asset system to remain usable in other platforms.

Why it matters: Interoperability reduces dependence on isolated applications and supports durable, machine-readable information. buildingSMART describes openBIM as improving accessibility, usability, management, and sustainability of digital built-environment data.

31. Industry Foundation Classes

Definition: Industry Foundation Classes, or IFC, is an open data standard for exchanging BIM and built-asset information.

AEC context: IFC can enable AI applications to analyze geometry, spaces, systems, components, relationships, and selected asset attributes across different software products.

Why it matters: Open, structured information can improve portability and reduce vendor dependence, although the quality of an IFC exchange still depends on how the source information was created and mapped.

32. Geographic Information Systems and Geospatial AI

Definition: A geographic information system, or GIS, manages, analyzes, and visualizes information connected to physical locations. Geospatial AI applies AI techniques to that information.

AEC context: Applications include site selection, utility analysis, environmental assessment, routing, infrastructure planning, climate-risk analysis, and portfolio management.

Why it matters: Buildings and infrastructure exist within geographic, environmental, demographic, and regulatory contexts. Combining AI, BIM, and GIS can improve planning and asset decisions.

33. Internet of Things and Sensor Data

Definition: The Internet of Things, or IoT, consists of connected physical devices that collect and exchange operational information.

AEC context: Sensors can report temperature, vibration, energy use, equipment condition, occupancy, air quality, location, moisture, and production activity.

Why it matters: Sensor data can support digital twins and predictive maintenance, but organizations must manage connectivity, calibration, cybersecurity, data ownership, and retention.

34. Edge AI

Definition: Edge AI performs AI processing on or near a device rather than transmitting all information to a centralized cloud platform.

AEC context: A site camera or piece of equipment might detect a safety condition locally and issue an alert without sending continuous high-resolution video off-site.

Why it matters: Edge AI can reduce latency, connectivity requirements, bandwidth use, and some privacy concerns. It also creates a distributed fleet of devices that must be secured and maintained.

35. Knowledge Graph

Definition: A knowledge graph represents entities and the relationships among them in a structured, machine-readable network.

AEC context: It may connect an owner, facility, project, contract, drawing, asset, space, equipment item, manufacturer, warranty, and maintenance requirement.

Why it matters: Knowledge graphs can help AI reason across disconnected systems and understand how project or asset information relates.

36. Ontology and Taxonomy

Definition: A taxonomy organizes information into categories. An ontology goes further by defining concepts, properties, and relationships within a domain.

AEC context: A firm might standardize how it names project types, disciplines, spaces, systems, risk events, change causes, assets, and project outcomes.

Why it matters: AI performs poorly when the organization uses several incompatible names for the same concept. Shared classifications make information easier to compare, integrate, and analyze.

IV. Data, Reliability, and Performance

37. Data Governance

Definition: Data governance establishes accountability, policies, standards, permissions, and controls for how information is created, classified, stored, accessed, shared, retained, and disposed of.

AEC context: It applies to project records, BIM data, client information, financial data, employee information, site imagery, sensor feeds, and intellectual property.

Why it matters: AI governance cannot succeed without data governance. Executives should know what data an AI system uses, who owns it, who may access it, and whether it may be used to train external models.

38. Data Quality

Definition: Data quality describes whether information is accurate, complete, consistent, current, relevant, and suitable for its intended purpose.

AEC context: Inaccurate cost codes, inconsistent activity names, incomplete BIM attributes, or outdated project records can produce unreliable AI analysis.

Why it matters: AI does not repair a weak information foundation by magic. It can automate the production of polished conclusions from flawed data, a particularly dangerous form of efficiency.

39. Data Lineage and Provenance

Definition: Data lineage records how information moves and changes across systems. Provenance identifies where information originated, who created it, and how it was processed.

AEC context: A decision-maker may need to trace an AI-generated forecast back to source schedules, cost reports, field records, assumptions, and model versions.

Why it matters: Lineage and provenance support trust, auditing, dispute resolution, regulatory compliance, and error investigation. NIST guidance specifically emphasizes documenting data sources, assumptions, limitations, and lineage.

40. Hallucination or Confabulation

Definition: Hallucination, also called confabulation, occurs when a generative AI system produces false, unsupported, or fabricated information that may appear credible.

AEC context: A system might invent a contract requirement, cite a nonexistent standard, misstate a specification, or fabricate a project fact.

Why it matters: Fluency is not evidence of accuracy. High-consequence outputs should be grounded in approved sources and independently verified. NIST identifies confabulation as a significant generative-AI risk.

41. Model Evaluation and Evals

Definition: Model evaluation, commonly called evals, is the systematic testing of an AI system's accuracy, reliability, safety, usefulness, and consistency.

AEC context: An estimating assistant might be tested against historical takeoffs, while a contract assistant might be tested using known clauses, exceptions, and deliberately difficult questions.

Why it matters: A successful demonstration is not proof of production readiness. Executives should require documented test cases, acceptance thresholds, failure analysis, and continuing monitoring.

42. Model Drift

Definition: Model drift occurs when an AI system's performance changes because real-world data, operating conditions, behaviors, or relationships differ from those present during development.

AEC context: A forecasting model trained on one market, project type, labor environment, or procurement model may become less accurate when those conditions change.

Why it matters: AI performance must be monitored after deployment. Approval at launch should not become permanent approval by inertia.

43. Human-in-the-Loop (Bookend)

Definition: Human-in-the-loop refers to a process in which a qualified person reviews, approves, corrects, or overrides an AI system's output or action.

AEC context: AI may prepare a code-analysis summary, payment recommendation, schedule intervention, safety alert, or contractual response, but an authorized professional makes the decision.

Why it matters: Human oversight is essential when outcomes affect health, safety, contractual rights, professional obligations, employment, or substantial financial commitments.

44. Explainability

Definition: Explainability is the extent to which people can understand how an AI system reached a conclusion or recommendation.

AEC context: A project executive should be able to understand which indicators caused a project to be classified as high risk.

Why it matters: Explanations help people challenge outputs, identify errors, satisfy oversight obligations, and decide whether an AI recommendation is reasonable.

45. Bias and Fairness

Definition: Bias is a systematic tendency that can create inaccurate or unequal outcomes. Fairness concerns whether an AI system's design and use produce unjustified differences in treatment or impact.

AEC context: Bias may affect recruiting, promotion, subcontractor evaluation, site selection, risk scoring, resource allocation, or safety monitoring.

Why it matters: Historical data may contain historical inequities. Automating those patterns can scale legal, reputational, workforce, and community risks.

V. Governance, Risk, and Security

46. Responsible AI

Definition: Responsible AI is the disciplined development and use of AI in ways that are lawful, safe, secure, transparent, accountable, privacy-conscious, and aligned with organizational values.

AEC context: Responsible use may require documenting the system's purpose, limitations, source data, review procedures, affected stakeholders, and escalation process.

Why it matters: Responsible AI turns broad ethical principles into operating requirements. It protects people and projects while enabling innovation to proceed with clear boundaries.

47. AI Governance and AI Management System

Definition: AI governance establishes authority, policies, roles, controls, decision rights, and accountability for AI. An AI management system organizes those activities into a repeatable enterprise framework.

AEC context: Governance determines who may approve tools, connect project data, authorize autonomous actions, accept risks, investigate incidents, and retire systems.

Why it matters: AI should be governed through established business, risk, information-security, legal, quality, and project-delivery structures. ISO/IEC 42001 provides a formal framework for managing AI risks and opportunities, while the NIST AI RMF organizes activities around Govern, Map, Measure, and Manage.

48. AI Risk and Impact Assessment

Definition: An AI risk or impact assessment evaluates a proposed system's purpose, affected parties, data, limitations, potential harms, controls, and residual risk before deployment.

AEC context: A low-risk writing assistant should not receive the same review as a system influencing hiring, safety, structural decisions, payment, claims, or autonomous equipment.

Why it matters: Risk-based review allows organizations to move quickly on low-consequence applications while applying stronger controls to consequential uses.

49. Guardrails

Definition: Guardrails are technical and procedural controls that limit what an AI system can access, generate, recommend, or do.

AEC context: Guardrails may prevent an assistant from accessing restricted projects, issuing external correspondence, modifying a schedule, approving a payment, or presenting an answer without source citations.

Why it matters: Guardrails convert policy into enforceable operating boundaries. They should be tested because instructions alone do not reliably constrain every AI behavior.

50. AI Security Threats

Definition: AI security threats include attempts to manipulate an AI system, steal information, corrupt its data, misuse its tools, or cause unauthorized actions. Examples include prompt injection, data poisoning, sensitive-information leakage, model theft, insecure integrations, and unauthorized "shadow AI" use.

AEC context: A malicious instruction hidden in a document could attempt to manipulate an AI assistant reviewing project files. Employees may also expose confidential drawings, contracts, or client information by entering them into unapproved tools.

Why it matters: AI expands the cybersecurity attack surface because models interact with natural language, documents, databases, software tools, and external content. OWASP identifies prompt injection, sensitive-information disclosure, supply-chain vulnerabilities, misinformation, and excessive agency among major generative-AI application risks.

Closing Perspective

AEC executives do not need to master every AI advancement. They do need enough fluency to distinguish a productivity tool from an autonomous decision system, a persuasive demonstration from a proven capability, and an experimental convenience from an enterprise risk. A shared glossary gives leadership teams the language to make those distinctions. Used properly, it becomes the front door to stronger investment discipline, better procurement, safer experimentation, clearer accountability, and more deliberate transformation.

Editorial and Source Note

These definitions are plain-language executive interpretations written for an AEC audience. They are not intended to replace technical standards, contractual definitions, legal advice, professional judgment, or project-specific requirements.

The glossary has been informed by terminology and management principles published by NIST, ISO, buildingSMART International, and the OWASP Generative AI Security Project.

This document was co-created by Digital Conditions and GPT-5.6 Sol LLM.