Construction AI Lab
Practical AI for better results
@ sudyco®

Construction AI Lab

# Technical Guidance for AI Integration

*Questions and minimum elements for leadership, IT, information security, legal, risk, records, HR, and operations teams supporting the Organizational AI Integration Standard.*

2027 Edition

**MASTER DRAFT v1**

## Construction AI Lab point of view

**Use AI to help people do better work. Do not let AI replace human judgment, responsibility, or accountability.**

This publication is intended as a practical, editable industry resource. Each organization should adapt it to its operations and have appropriate legal, information security, privacy, human resources, risk, records-management, and contractual review before adoption.

# Purpose of This Guidance

This document does not prescribe an IT architecture or claim to replace technical, legal, cybersecurity, privacy, or compliance expertise. It gives construction organizations a practical set of questions and minimum governance elements to address when AI enters the business through standalone tools, enterprise platforms, project systems, embedded features, integrations, agents, and vendor services.

### The relationship between the two documents

The Organizational AI Integration Standard states what the organization is trying to protect and achieve. This Technical Guidance helps leadership and technical functions create the controls, decisions, and records needed to support that Standard.

# Who Should Use It

AI integration should not be assigned to IT alone. Use this guidance in a cross-functional working session that includes, at minimum:

- Executive sponsor
- Operations or project leadership
- IT and information security
- Legal, contracts, and risk
- HR and training
- Records and information management
- Safety and quality, when relevant
- People who actually perform the affected work

# How to Use the Questions

1. List the AI platforms, embedded features, integrations, agents, and use cases currently in use or under consideration.
2. Answer each section at a level appropriate to the risk and scale of the use.
3. Record decisions, owners, required controls, open issues, and review dates.
4. Do not enable or expand higher-risk uses until accountable leaders understand and accept the residual risk.
5. Revisit answers when vendors, features, data practices, contracts, laws, or actual experience change.

# 1. Inventory and Ownership

Know what is entering the organization and who is accountable.

**Questions to answer**

- What AI systems, embedded features, APIs, plug-ins, agents, and automations are being used, tested, or requested?
- Which uses are official, experimental, personal, shadow, or vendor-enabled by default?
- Who owns each platform and use case from a business, technical, security, and risk perspective?
- Who has authority to enable, restrict, suspend, or retire an AI capability?
- Where is the current inventory maintained, and how often is it reconciled against actual use?

**Minimum documented elements**

- AI system and use-case inventory

- Named business owner and technical owner
- Enable/disable authority
- Review date and change trigger

## 2. Approved Platforms and Vendor Due Diligence

Understand the vendor and the boundaries of its responsibility.

**Questions to answer**

- Is the product approved for the intended data, users, projects, and decisions?
- What model providers or subprocessors are used, and can they change without notice?
- Is customer data used to train, improve, or evaluate models? What contractual terms apply?
- What retention, deletion, residency, backup, portability, and termination provisions apply?
- What security, privacy, assurance, audit, incident-notification, availability, and support commitments exist?
- What AI features can the vendor enable by default, and how will changes be reviewed?
- What responsibilities remain with the construction organization even when the vendor provides safeguards?

**Minimum documented elements**

- Approved vendor list
- Documented due diligence
- Contract and data-use review
- Vendor-change monitoring

## 3. Data Classification and Use Rules

Decide what information may enter which systems.

**Questions to answer**

- What categories of company and project data exist, and which are permitted, restricted, or prohibited in each AI system?
- How are owner, designer, subcontractor, employee, applicant, public, security-sensitive, privileged, personal, export-controlled, and regulated data handled?
- Do contracts or agency requirements limit AI processing, storage, location, disclosure, or reuse?
- Can users recognize protected information before entering it?
- How will prompts, uploaded files, conversation history, generated outputs, and system logs be retained or deleted?
- What process exists for correcting, removing, or responding to unintended disclosure?

**Minimum documented elements**

- Data classification mapped to AI systems
- Prohibited-data rules
- Retention and deletion decisions
- Contract-specific restrictions

## 4. Identity, Access, and Permissions

Ensure AI cannot reach more than the user and organization intend.

**Questions to answer**

- Is enterprise identity, multifactor authentication, role-based access, and least privilege used?

- Do AI features inherit the underlying platform permissions accurately?
- Can the system search, summarize, create, modify, publish, send, approve, or execute actions?
- Are service accounts, integrations, agents, and delegated permissions separately controlled?
- How are project transfers, role changes, terminations, joint ventures, consultants, and temporary users handled?
- Can sensitive actions require a second person, additional authentication, or explicit approval?

**Minimum documented elements**

- Role and permission model
- Privileged-action controls
- Joiner/mover/leaver process
- Agent and service-account controls

## 5. Human Review and Action Boundaries

Convert the organizational principle into system and workflow controls.

**Questions to answer**

- Which outputs may be used as drafts, recommendations, summaries, or research aids?
- Which actions must always remain subject to qualified human review and approval?
- Which decisions or transactions must never be delegated to AI?
- How will the system make draft status, sources, uncertainty, and required review visible?
- Can automated workflows bypass normal approvals, separation of duties, or contractual authority?
- What evidence will show who reviewed, changed, approved, and released the result?

**Minimum documented elements**

- Use-case decision boundaries
- Required reviewer qualifications
- Approval and audit trail
- Prohibited autonomous actions

## 6. Accuracy, Testing, and Validation

Test the actual use, not only the vendor's general claims.

**Questions to answer**

- What could go wrong in this use case, and what would the consequence be?
- What representative project data, edge cases, adversarial cases, and failure conditions will be tested?
- How will factual accuracy, completeness, source fidelity, bias, reliability, repeatability, and usability be evaluated?
- Does testing include integrations, permissions, internet access, external actions, and behavior outside the intended sandbox?
- What performance level is acceptable, and who decides?
- How will changes in models, prompts, configurations, data, or workflows trigger retesting?

**Minimum documented elements**

- Use-case test plan
- Acceptance criteria
- Qualified reviewers
- Retest triggers

## 7. Security, Monitoring, and Incident Response

Prepare for misuse, unexpected behavior, and compromised systems.

**Questions to answer**

- What threats arise from prompt injection, malicious files, data poisoning, excessive permissions, insecure integrations, compromised accounts, or manipulated outputs?
- What activity is logged, who reviews it, and how long is it retained?
- Can anomalous access, bulk extraction, unauthorized actions, or policy violations be detected?
- How can an AI feature, integration, credential, or agent be quickly contained or disabled?
- Does the incident process address AI-specific evidence, vendor coordination, affected parties, legal holds, contractual notice, and lessons learned?
- Are users told how and where to report suspicious behavior or material errors?

**Minimum documented elements**

- AI threat review
- Logging and monitoring plan
- Containment capability
- AI incident procedure

## 8. Records, Transparency, and Traceability

Preserve the record needed to understand important work.

**Questions to answer**

- When is AI involvement material enough to disclose or record?
- Which prompts, sources, outputs, edits, approvals, and system versions must be retained?
- Can a reviewer trace an important output back to authoritative source material?
- How will AI-assisted records be classified, searched, produced, corrected, and disposed of?
- What disclosures are needed for owners, regulators, courts, auditors, professional reviewers, employees, or the public?
- Can the organization explain the role AI played without overstating what the system knows or decided?

**Minimum documented elements**

- Disclosure rules
- Traceability requirements
- Records schedule
- Source and approval evidence

## 9. Training, Support, and Practical Usability

Make the approved path easier than the shadow path.

**Questions to answer**

- Do users understand approved tools, permitted data, review duties, limitations, and reporting expectations?
- Is training specific to construction roles and actual workflows rather than generic AI awareness?
- Can users get timely help when they are unsure?
- Are approved tools practical enough that users will not bypass them?
- How will the organization identify overreliance, skill erosion, confusion, or declining collaboration?
- How will lessons from projects and users be shared without exposing protected information?

**Minimum documented elements**

- Role-based training
- Support channel
- User feedback process
- Capability and overreliance monitoring

## 10. Ongoing Review and Change Management

Treat AI integration as a living management responsibility.

**Questions to answer**

- Who reviews vendor releases, model changes, new AI functions, contractual changes, and regulatory developments?
- What changes require notice, approval, retraining, retesting, or revised procedures?
- How are benefits and unintended consequences measured together?
- What evidence would cause the organization to pause, limit, redesign, or retire a use?
- How often will leadership review the AI inventory, incidents, performance, lessons, and Early Warning Signs?
- How will the annual Organizational AI Integration Standard review be coordinated with technical updates?

**Minimum documented elements**

- Change-review process
- Annual leadership review
- Benefit-and-risk measures
- Pause/retire criteria

## Cross-Functional Decision Record

Use the following record for each significant AI platform, embedded feature, integration, agent, or use case. Expand it when the risk requires more detail.

| | |
|---|---|
| **System / feature / use case** | [Complete] |
| **Business purpose and expected value** | [Complete] |
| **Affected workflows and people** | [Complete] |
| **Data involved and restrictions** | [Complete] |
| **Human review and decision boundary** | [Complete] |
| **Key risks and required controls** | [Complete] |
| **Vendor / technical owner** | [Complete] |
| **Business owner / accountable executive** | [Complete] |

| | |
|---|---|
| **Approval, conditions, or rejection** | [Complete] |
| **Review date and change triggers** | [Complete] |

## Leadership Review Questions

**Before approving a significant AI use**

Leadership should be able to answer these questions in plain language - not only technical language.

6. What work will this AI help people do better?
7. What could it weaken - including trust, transparency, collaboration, communication, judgment, responsibility, accountability, safety, quality, or skills?
8. Who remains accountable, and what human review is required?
9. What information can the system access, retain, disclose, or act upon?
10. What safeguards are provided by the vendor, and what responsibilities remain ours?
11. How will we know whether it is working, drifting, or causing unintended consequences?
12. Can we pause, contain, correct, or retire it when needed?
13. Is the organization prepared to absorb this change responsibly?

## Reference Basis and Important Notice

This guidance is designed to support, not replace, the work of qualified IT, information security, privacy, legal, risk, HR, records-management, safety, quality, and other professionals. Organizations seeking a formal AI management system or certification should evaluate applicable standards and legal requirements with qualified advisers.

**Reference sources:** NIST AI Risk Management Framework 1.0; NIST Generative AI Profile (NIST AI 600-1); NIST AI Resource Center materials; ISO/IEC 42001:2023 public overview; and public Procore guidance on permission-based AI access and human-in-the-loop verification. Accessed August 2026.