



FOR ENGINEERING AND CONSTRUCTION COMPANIES

AI Use Policy for Contractors

A policy a construction company can adopt in a week. Replace every [bracket], have counsel read it once, and it is yours. Written to satisfy what carriers, owners and counsel now ask for: a named owner, an approved-tool list, a human on every output, confidential data kept inside the company, and a record you can produce.

1. Purpose and scope

[Company name] uses AI tools to do work faster and better. This policy sets the rules so that use never puts a project, a client, an employee or the company at risk.

It applies to every employee, contractor and temporary worker of [Company name], on every device, for any AI tool that touches company work: drafting, summarizing, estimating, scheduling, design, correspondence, RFIs, submittals, change orders, claims, safety, HR and finance.

It covers three kinds of tool:

- Public tools anyone can open in a browser (consumer chat assistants, free image and transcription tools).
- Enterprise tools the company licenses, with a contract that says what happens to our data.
- AI built into software we already run (estimating, project management, document control, email).

The policy does not change who is responsible for the work. Output from an AI tool is treated exactly like work an employee did by hand: the person who uses it owns it.

2. Definitions

- **AI tool.** Any software that generates, summarizes, classifies, predicts or decides using a machine-learning model, whether standalone or built into another product.
 - **Approved tool.** An AI tool on the inventory in Section 4, used within the limits recorded there.
 - **Public tool.** An AI tool without a company contract governing our data. Assume anything entered into it may be stored, read by the vendor, used to train models, or produced in a dispute.
 - **Confidential information.** Anything not already public: bids, estimates, unit rates, subcontractor pricing, contracts, claims, schedules, drawings, safety records, personnel and financial data, and client and owner information of every kind.
 - **Human review.** A named, qualified person reads the AI output, checks it against the source it should rest on, and takes responsibility for it before it is relied on, sent or filed.
-

3. Governance

[Company name] has an AI Committee. It owns this policy and every decision under it.

SEAT	WHO	WHAT THEY DECIDE
Chair	[Name, title]	Final say on the policy and on exceptions
Legal	[Name or outside counsel]	Confidentiality, privilege, contract and liability questions
IT / Security	[Name]	Technical vetting, access, logging, blocking
Operations	[Name, estimating or project executive]	Which uses are worth the risk, and what training the field needs
Risk / Insurance	[Name or broker]	Coverage questions and incident review

The committee meets at least [Quarterly] and whenever a new tool or a new use is proposed. It keeps minutes. Its decisions are published to staff within [5] business days.

No AI tool is bought, trialed on company data, or connected to a company system without the committee's written approval. A manager cannot approve one alone.

4. Approved tools

IT keeps the inventory. If a tool is not on it, it is not approved. The inventory is visible to every employee, so nobody has to guess.

TOOL	VENDOR	APPROVED USES	DATA ALLOWED	OWNER	APPROVED ON	REVIEW BY
[Tool]	[Vendor]	[First drafts of correspondence]	[Internal, non-confidential]	[Name]	[Date]	[Date]
[Tool]	[Vendor]	[Estimating and contract analysis on live projects]	[Confidential, inside our environment]	[Name]	[Date]	[Date]

To add a tool, the requester sends the committee a one-page request: what it does, what data it will see, where that data goes, and who will use it. IT completes the vendor check in Section 9. The committee approves, limits or declines in writing.

Public tools are blocked on company devices and networks unless the committee has listed one with an explicit data limit. Personal devices used for company work follow the same list.

5. Acceptable and prohibited uses

Acceptable, with an approved tool and human review

- First drafts of letters, emails, meeting summaries and internal reports.
- Summarizing documents the company already holds, for the people entitled to read them.
- Research and background on codes, standards, products and methods, checked against the primary source before use.
- Estimating, scheduling, contract analysis and project controls, only in a tool approved for confidential data and only where the output shows its sources.
- Transcribing and summarizing company meetings, with every attendee told a recorder is on.

Prohibited

- Entering confidential information into a public tool. No exceptions.
- Sending, filing or relying on AI output nobody has reviewed.
- Using AI output as a design, engineering or safety determination without a licensed or qualified person's review and sign-off.
- Asking an AI tool for legal advice on a live matter instead of counsel.
- Presenting AI-generated text, images or figures as a person's own work to a client, owner, agency or court without saying so where the recipient would expect to know.
- Using AI to make or recommend employment decisions (hiring, discipline, termination) without HR and counsel review.
- Recording or transcribing anyone without the notice the law and the company require.

Anything not listed is referred to the committee before it is done.

6. Confidential information, privilege and discoverability

Three facts drive this section.

1. **What goes into a public tool may not come back.** Vendors can store it, read it, and use it to train models. A trade secret entered into a public tool can lose its protection. Subcontractor pricing entered into one can end up in front of a competitor.
2. **Privilege is fragile.** Communications with counsel are protected; a question typed into a chatbot is not, and routing a legal question through a tool can waive the protection the lawyer's answer would have carried. Legal questions go to counsel.
3. **Everything can be discovered.** In a claim or a lawsuit, prompts, outputs and logs can be demanded, including from closed, enterprise systems. Assume anything typed into any AI tool may be read aloud in a deposition.

So: confidential information goes only into tools the committee has approved for it, inside an environment the company controls, with a record of who asked what and what came back. Staff write prompts the way they write emails: as if the other side will read them one day.

The other thing a tool takes: how your people decide

Consent agreements and data processing terms classify data: the file you upload, the text you type, "customer content." Most now promise that content is not used to train models. Few say the same about everything around it: which questions your people ask, in what order, what they accept, what they reject, what they rewrite, and how long they weigh each answer. Vendors file that under usage data, interaction data, telemetry or service improvement, reserve it under a different clause, often a different consent, and often for themselves.

For a contractor, that is the company's judgment, captured one prompt at a time. How your chief estimator decides what to carry. Which clauses a project executive escalates and which he lets go. What your best superintendent flags on a walk. None of it sits in a file, so none of it is "customer content," so the file-based protection in the agreement never reaches it. Over months, a frontier model app can hold a working model of how your company makes decisions, and the consent you signed never named it.

This policy treats decision patterns as confidential information. Before any AI tool is approved for work that involves judgment (estimating, contracts, scheduling, bid or no-bid, claims), the committee confirms in the contract, not on the privacy page:

- Interaction and usage data is not used to train or improve models for anyone else.
- It is not retained beyond what operating the service for us requires.
- The company can see it and export it.
- Where the pattern of our decisions is kept at all, it is kept as our own record, in an environment we control, and it leaves with us.

7. Human review and accountability

"The AI did it" is not a defense, inside the company or outside it. The company is responsible for AI output as if an employee wrote it by hand, and the employee who used the tool is responsible for checking it.

Review means more than reading. The reviewer:

- Checks every figure, date, clause reference and citation against the source it claims to rest on. A number with no source is treated as wrong until shown otherwise.
- Confirms the output answers the actual question on the actual project, not a similar one.
- Signs or initials the output, or records the review in the tool, before it is sent, filed or acted on.

The higher the stakes, the more senior the reviewer. Minimum review levels:

OUTPUT	MINIMUM REVIEWER
Internal draft, meeting summary	The person who ran the tool
Correspondence to a client, owner, sub or agency	The person's manager
Estimate, bid figure, schedule, cost report	The responsible estimator, scheduler or PM, by name
Contract interpretation, notice, change order, claim	Project executive and, where the policy or the contract requires, counsel
Design, engineering or safety determination	The licensed or qualified person who will sign it

An employee who cannot explain how an output was reached, and where its facts came from, does not send it.

8. Records

For every approved tool that handles confidential information, the company keeps, for at least the retention period that applies to the project record it touched:

- Who asked, when, and on which project.
- What was asked and what came back.
- Which sources the output drew on.
- Who reviewed it, and what they changed or rejected.
- When the tool declined to answer, and why.

These records live where the project record lives, under the same access controls, not in a personal account or a vendor's console the company cannot export.

The records exist for two reasons. First, so that a figure in a bid or a claim can be traced back to its inputs years later. Second, so that when a carrier, an owner, an auditor or a court asks how an output was produced, the company can show them instead of telling them.

9. Vendors

Before a tool goes on the inventory, IT answers these in writing and the committee reads the answers. A vendor who cannot answer in plain terms is not approved.

QUESTION	ANSWER NEEDED
Where is our data stored, and in which country?	Named region; nothing outside the regions the company allows
Is our data used to train models, for us or for anyone else?	No, in the contract, not in a help page
Is our data mixed with other customers' data?	No
Can we export everything and have it deleted when we leave?	Yes, with a timeline
Who at the vendor can see our data, and when?	Named roles, logged access
Does the tool show the sources behind each output?	Yes for any tool that touches figures, contracts or schedules
Can we see a log of every request and every answer?	Yes, exportable
Can we control which models it uses and pin a version?	Yes, or the vendor explains why not
What independent security audit has it passed?	Report in hand, dated
What happens when the tool is unsure or should not answer?	It says so and records it, rather than guessing
Is how our people use the tool (prompts, choices, corrections, sequence) treated as our confidential data, or as the vendor's usage data?	Ours, in the contract; not used to train or improve models for others; exportable; kept in our environment

The contract names the data limits from the inventory. If the vendor changes terms, the tool goes back to the committee.

10. Training

Everyone covered by this policy completes AI training within [30] days of hire and every [12] months after, with a short assessment. Nobody gets access to an approved tool before completing it.

Training is by role, because the risks are different:

- **Everyone:** what a public tool does with what you type, what confidential means here, how to recognize an invented fact or citation, and how to report one.
- **Estimators, schedulers, PMs, contracts:** checking output against sources, the review levels in Section 7, and what the record in Section 8 must show.
- **Field supervision:** photos, voice notes and transcripts as project records; consent before recording; what not to say to a tool.
- **Managers:** the approval process, the inventory, and how to handle an employee's request to use something new.

- **Committee:** vendor questions, insurance developments and incident review.

Leadership takes the same training first, and uses the approved tools in the open. Culture outlasts the document: people do what they see their managers do.

11. Monitoring, incidents and discipline

IT reviews access and usage logs for approved tools at least [Monthly] and reports to the committee: who is using what, on which projects, and any attempt to reach a blocked tool.

An **incident** is any of the following, and is reported to [Committee chair] the same day it is found:

- Confidential information entered into a public or unapproved tool.
- AI output sent, filed or relied on without review, or found to be wrong after it was.
- An invented fact, figure, clause or citation that reached a client, owner, sub, agency or court.
- A vendor breach, outage or change of terms affecting company data.
- A recording made without the required notice.

The committee logs every incident, decides what to tell the affected party, the carrier and counsel, and records what changed so it does not happen again. Reporting a mistake you made yourself is never, by itself, a disciplinary matter. Hiding one is.

Violations of this policy are handled under [Company name]'s standard disciplinary process, up to and including termination. Deliberately entering confidential information into a public tool, or sending unreviewed output to a client, is treated as a serious violation on the first occurrence.

12. Insurance

Coverage for AI-related claims is changing, and not in the policyholder's favor. Endorsements that exclude losses arising from generative AI are now being offered on general liability forms, and how professional liability policies treat an AI-assisted error is unsettled. Carriers are also beginning to ask, at renewal, whether the company has an AI policy at all.

[Risk lead name] reviews this with the broker before every renewal and whenever a new tool is approved, and puts the answers in writing:

1. Does any policy we hold carry, or is any carrier proposing, an exclusion for claims arising from generative AI? On which forms?
2. If an estimate, schedule, design input or contract notice produced with an approved tool turns out to be wrong, which policy responds, and what would the carrier need to see?
3. Does the carrier require an AI policy, training records or tool inventory as a condition of coverage or pricing?
4. Does our cyber policy cover a vendor breach of data we entered into an approved tool?
5. What records would the carrier want at the time of a claim, and do we keep them (Section 8)?

The answers go in the committee's minutes. Where a gap is found, the committee decides whether to change the inventory, the use, or the coverage.

13. Ownership, review and acknowledgement

This policy is owned by [Committee chair]. It is reviewed every [6] months, and sooner when the law, a carrier, a client contract or an incident calls for it. Each version is dated and the previous one kept.

Every person covered signs an acknowledgement on hire and after each revision. The signed copy is kept with their personnel file.

ITEM

VALUE

Policy version

[1.0]

ITEM	VALUE
Effective	[Date]
Owner	[Name, title]
Next review	[Date]
Approved by	[Name, title, signature]

Acknowledgement. I have read [Company name]'s AI Use Policy. I understand which tools are approved, what I may not enter into any tool, that I am responsible for reviewing any AI output I use, and how to report an incident.

Name: ____ Signature: ____ Date: ____

Appendix: where a company running SymphonIQ already satisfies each section

The policy above is vendor-neutral and stands on its own. This page is for companies that run SymphonIQ, or are deciding whether to. It maps each requirement to what the platform does by design, so the committee can tick the box instead of building the control.

POLICY REQUIREMENT	HOW SYMPHONIQ MEETS IT
4. Approved-tool inventory, with data limits	One platform, one inventory entry. Every calculation tool and skill inside it is registered and logged; nothing runs that is not on the list.
4. Public tools blocked	Staff use SymphonIQ's own app, or Claude or ChatGPT as a front door into the same private environment. Nothing leaves it.
5. Confidential data only where output shows its sources	Every answer carries its receipts: the clause, the figure's inputs, the document and page it came from.
6. Nothing pooled, nothing trained across customers	One private environment per customer, in your own cloud account if you choose. Your memory stays yours.
6. Decision patterns stay the company's	What your experts teach it stays yours. Memory Fabric holds every project, and every decision on it, as your record in your environment. Nothing pooled, nothing trained across customers, and that includes interaction data. What one job teaches is organized for your next job and no one else's.
6. Discoverable record you would want discovered	The record shows who asked, what came back, which sources, who signed. It is built to be produced.
7. Human review before anything counts	Governance applied before it reasons, checks while it works, review before anything counts. An authorized person confirms; the record shows who.
7. Figures a reviewer can check	Facts are forced through deterministic calculation tools (200+). Same inputs, same answer, every time. The model narrates; the tools decide.
8. Records kept with the project record	Audit trail to source on every answer, kept in your environment under your access controls, exportable.
9. Model control and version pinning	Your IT team controls the models. 13 providers available, each pinned to a version you approve.
9. Tool says when it should not answer	A refusal log: every time a specialist declined to answer, and why, is recorded and reviewable.
9. Independent security audit	SOC 2 Type I audit underway, Type II to follow.
11. Usage monitoring	Every call logged, by person, project and specialist, visible to your administrators.

12. What a carrier would want to see	The signed, sourced, computed record above is the evidence a carrier, an auditor or a court asks for. SymphonIQ makes no representation about coverage; take this appendix to your broker.
--------------------------------------	--

Questions, or a copy of this policy in your own template: hello@symphoniq.ai

SymphonIQ, Inc. · symphoniq.ai · Version 1.1.5 · October 2026. Free to copy, adapt and distribute inside your company. This document is a template and is not legal advice.