AI Governance Guide

How The kAIzen Alliance LLC Uses AI in Client Work

Version 2.0 · July 2026 · The kAIzen Alliance LLC

The kAIzen Alliance LLC uses AI carefully, transparently, and with human judgment in charge. Responsible AI is not a separate policy we point to after the work is done. It is part of how we design, build, review, and hand off every client engagement. This guide explains, in plain language, how we handle your information, where we keep people accountable, and the standards we hold ourselves to. It is written to be shared with the people who need to review it, including procurement, legal, and internal stakeholders.

> **The short version.** AI supports decisions. People own them. We keep sensitive information out of AI tools, use vetted paid-tier tools for client work, and require human review before anything is delivered or used to make a decision.

Our Responsible AI Principles

These principles guide how we design, review, and hand off AI-supported work.

Humans Remain Accountable. AI supports decisions. It does not make them. Every system has a clear human owner for review, approval, and override.

No Black Boxes. You should understand what the system does, where its limits are, and when human review is required.

Authenticity Is Protected. AI can mimic empathy. It cannot feel it. We protect the human judgment, care, and trust that should not be automated.

Risk Awareness Is Built In. Hallucinations, data leakage, and bias are design issues, not afterthoughts. We plan for them before the work goes live.

Capability, Not Dependency. The goal is for you to own and operate the capability independently. We build organizational strength, not reliance on us.

Documentation Matters. Important AI-related decisions are documented so they can be reviewed, explained, and improved later.

What This Means in Client Work

Responsible AI shows up in ordinary decisions, not just formal policy language. Before we use AI in an engagement, we look at what kind of information is involved, what the output will influence, and how much review the work needs before anyone relies on it.

Some work is low risk: brainstorming from public information, summarizing a published source, or drafting early options that will be revised before use. Some work needs more control: client-specific strategy, internal operations, financial data, or anything that affects decisions about people, money, contracts, or reputation.

The higher the risk, the more review we apply. Human judgment stays in charge at every level.

Client Data Boundaries

Not every kind of information belongs in an AI tool. We draw clear lines before the work begins.

Sensitive Information Stays Out. Personally identifiable information, financial account details, credentials, regulated information, and trade secrets do not go into AI tools.

Redaction Comes First. Before files are used in AI-assisted work, sensitive information is removed or replaced. When testing is needed, we can create sample data that reflects the workflow without exposing real people or private details.

Paid Tools Only for Client Work. We use vetted paid-tier tools for client work, because business information needs stronger privacy and no-training protections than free tools usually provide.

You Can Set Tool Boundaries. If you do not want a specific AI tool used in your engagement, tell us at the start and we will work around it where the project allows.

Common Questions

Will our data be used to train AI models?

We choose paid-tier tools with stronger privacy and no-training commitments for client work. We do not intentionally use tools that train public models on your business information.

Does a human review AI output?

Yes. AI output is a starting point, not the final answer. Human review is required before anything is delivered, published, or used to support a decision.

Can AI make decisions for our business?

No. AI can help organize information, draft options, analyze patterns, or support a recommendation. People remain responsible for decisions, approvals, and final judgment.

What information should never go into AI?

Personally identifiable information, financial account details, passwords, API keys, regulated information, and trade secrets should not be entered into AI tools.

Can we ask you not to use a specific AI tool?

Yes. If you have concerns about a particular tool or vendor, tell us at the start of the engagement and we will work around it where the project allows.

What happens if AI makes a mistake?

We treat mistakes according to their seriousness. Minor issues are corrected internally. If an error affects your work, your information, or a material decision, we stop, contain the issue, tell you what happened, and fix the process that allowed it.

Talk to Us

If your team needs a deeper review for procurement, legal, or internal approval, we are glad to walk through any of this directly. The kAIzen Alliance LLC, Florence, Arizona. Start a conversation at thekaizenalliance.com/contact.