adaptcx

BUYER GUIDE

The AI Vendor RFP Checklist

What to ask before you buy AI for regulated CX - 47 questions organized by what they actually tell you.

A practical evaluation framework for procurement, technology, security and operations leaders. AdaptCX builds on an earlier, smaller version. The expanded platform has not yet been deployed in full at a customer site.

GUIDE MAP

Contents

Why most AI RFPs miss the point for the failure modes that matter Section D - Implementation and the yet

SECTION

Why most AI RFPs miss the point

Most procurement teams write AI RFPs the same way they write RFPs for traditional contact-center software: feature checklists, pricing matrices, integration questions, security questionnaires. This produces RFP responses that are technically correct and operationally useless. The reason: AI products fail differently than traditional software. They fail in production, at scale, in ways that look right in a demo and wrong six months in. Things like inconsistent answers to the same question, hallucinated policy citations, model drift after a vendor update, audit trails that don't hold up to a regulator's review. None of these failure modes appear in a feature checklist. This playbook gives you the 47 questions that actually predict production success. They are organized by what they tell you - not by which vendor team they come from If a vendor can't answer the questions in Section 03 (Reliability and Failure) directly, treat it as a red flag, not a small one. Section A - Architecture and deployment These questions tell you whether the vendor's product can survive in your environment without re-authorizing your security boundary - Is the deployment fully in-tenant on a cloud we already operate, or vendor-hosted with our data crossing your perimeter? - If in-tenant, who owns the IAM, the encryption keys, the audit sinks, and the

SECTION

Why most AI RFPs miss the point (continued)

network policies? - Walk me through the data flow for a single customer interaction, naming every cloud service touched and which side of the trust boundary it's on. - What does a deployment in our existing security boundary look like - what artifacts are shared, what access is required, and what stays inside our perimeter? - Can you point to a customer with a similar compliance posture to ours that you have already deployed to in production? - If our security team finds an issue post-deployment, what's the patch and re-validation process? Section B - Reliability and failure modes

SECTION

for the failure modes that matter

AI systems fail in characteristic ways that traditional software does not. These questions probe - What happens when your AI gives the wrong answer? Walk me through your incident process from detection to resolution. - How do you ensure the same question gets the same answer across calls - at scale, across agents, over time? - Show me the audit log for a single interaction: every retrieval, every model call, every override. Is that log tamper-evident, and who has write access to it? - How do you handle model updates from your underlying LLM provider? What's the testing and rollout process before changes hit production? - What's your SLO for response latency at p50, p95, p99? Do those SLOs hold under realistic concurrent load? - What happens when the AI is uncertain - does it say so, route to a human, or guess confidently? Show me an example. - If we needed to roll back to last quarter's behavior because a model update degraded responses, can we? The honest answer to several of these is 'we're working on it.' The dishonest answer is to claim full coverage. Note which vendors do which. Section C - Compliance posture (without the marketing) Compliance language in AI RFP responses is often performative. These questions force specifics - What compliance frameworks are you certified against, by whom, with what audit date and

SECTION

for the failure modes that matter (continued)

expiration? - If you describe yourself as 'aligned' or 'aware' rather than 'certified', explain the gap. Specifically - what controls would need to be added for full certification? - Who is the controller and who is the processor in our deployment? Show me the data processing agreement. - Where is customer data stored at rest, where is it processed in transit, and where do logs and telemetry land? - Do model providers (OpenAI, Anthropic, Google) see our data? If yes, under what terms? If no, how do you architecturally prevent it? - What is your incident notification timeline if you become aware of a breach affecting our data?

SECTION

Section D - Implementation and the

- Have you ever had a security incident affecting customer data? If yes, walk me through it. first 90 days Most AI projects fail not in the architecture but in the implementation. These questions surface that risk - What's the realistic timeline from contract signature to first production traffic - your timeline, not the marketing timeline? - Who from your team is on the implementation? Are they full-time on us or shared across customers? - What does a successful first 90 days look like - what milestones, what metrics, what definition of done? - What goes wrong most often in the first 90 days, and how do you handle it? - What are we responsible for and what are you responsible for? Be specific - list ten implementation tasks and assign owners. - If we want to bring our own knowledge base, what does data preparation look like? Quote effort in person-weeks, not 'it depends'. Section E - Ongoing operations and contract structure The contract sets the next three years of vendor relationship. Most teams under-negotiate it - What's the unit of pricing - per user, per interaction, per token, per month? Show me a sample bill. - How does the price scale at 2× and 5× our current volume? - What are the termination clauses - for cause, for convenience, for cause-by-you? - If the contract ends, what happens to our data, our trained models,

SECTION

Section D - Implementation and the (continued)

our knowledge base, our audit logs? What's the export format? - What's your renewal price escalation policy? Is it tied to an index or capped? - Do you have an exit assistance provision - and what does it actually include? - Who do we call when something breaks - name an actual person, not 'support'? How to use this checklist Print this list. Hand it to every vendor in your shortlist. Tell them you want written answers, not a sales call. Then read the answers carefully. The good vendors will write a careful, honest response that admits uncertainty in a few places. The bad vendors will write marketing copy

SECTION

yet

that doesn't actually answer the question. Score each response on a simple 0-2 scale: 0 = didn't answer or answered with marketing language, 1 = partially answered, 2 = answered specifically and honestly. The vendor with the highest score after honest evaluation is almost always the right vendor. If your top three vendors all score below 70%, that's a sign your RFP needs to push harder - or that the AI category for your specific use case isn't mature enough

NEXT STEP

Continue the conversation

If this guide surfaced questions about your architecture, compliance posture, deployment model or AI operating model, talk with an AdaptCX engineer or solutions architect.

Talk to AdaptCX →