

AI VENDOR DUE DILIGENCE ESSENTIAL QUESTIONS

Ethicore Advisors | ethicoreadvisors.com

Vendors who can't or won't answer these questions are ones you should avoid.

1 Training Data Sourcing

- 1 Where does your training data come from?**
→ Expect specific sources, not vague 'publicly available data'
- 2 What licenses or permissions do you have for training data?**
→ Understand the legal basis for training data use
- 3 How do you verify training data doesn't include copyrighted, private, or harmful content?**
→ Data quality and filtering processes
- 4 Can you provide documentation of training data sources and licenses?**
→ Vendors with good practices document this
- 5 How often do you update or retrain your models, and with what data?**
→ Ongoing data governance, not just initial training

2 IP Indemnification

- 6 Do you provide IP indemnification for AI-generated outputs?**
→ Full indemnification is ideal but rare - understand your exposure
- 7 What are the limits on your IP indemnification?**
→ Caps, exclusions, and conditions all matter
- 8 What happens if your AI generates content that infringes copyright or trademark?**
→ Know whether you're on your own for legal defense
- 9 Do you have IP insurance covering AI-generated content?**
→ An additional protection layer if the vendor has meaningful coverage
- 10 Have you faced IP infringement claims related to your AI?**
→ Past claims indicate higher risk

3 Bias and Fairness

- 11 How do you test for bias in your AI systems?**
→ Look for specific methodologies, not just "we test for bias"
- 12 Can you provide the results of your bias testing?**
→ Transparent vendors will share this
- 13 What demographic groups are included in your testing data?**
→ Limited diversity in testing creates blind spots
- 14 How do you monitor for bias in production?**
→ Ongoing monitoring catches emerging bias
- 15 What do you do when bias is discovered?**
→ Understand the timeframe for fixes

4 Privacy and Data Governance

- 16 What data from my organization does your AI access or process?**
→ Identify potential privacy issues clearly
- 17 Where is data stored and processed geographically?**
→ Matters for GDPR and other regulatory compliance
- 18 How do you secure data processed by your AI?**
→ Encryption, access controls, and audit logs

19	Can data be deleted on request? → Understand deletion processes and timelines
20	Do you use customer data to train or improve your AI? → Potentially a significant privacy concern - understand opt-out
21	Who has access to data processed by your system? → Vendor employees, subcontractors, and other customers?

5 Sustainability and Environmental Impact

22	What is the energy consumption of your AI system? → Any metric helps - per query, per user, per hour
23	What is the carbon footprint of your AI operations? → Many vendors don't calculate this, but asking prompts consideration
24	Do you use renewable energy for AI infrastructure? → Understand their renewable energy commitments
25	How do you optimize for energy efficiency? → Efficiency improvements reduce both costs and environmental impact
26	Can you provide environmental impact reporting? → Needed for ESG compliance

6 Vulnerable Population Protections

27	How does your AI handle interactions with children? → COPPA compliance and child safety considerations
28	How does your AI accommodate users with disabilities? → Accessibility requirements matter
29	Can users easily escalate from AI to human assistance? → Critical for vulnerable populations
30	How do you handle users who may be experiencing a crisis? → Mental health, financial distress, and other vulnerable states
31	Do you provide training on considerations for vulnerable populations? → Indicates vendor sophistication

7 Transparency and Explainability

32	Can you explain how your AI makes decisions? → Black-box systems create accountability problems
33	What transparency features do you provide to end users? → Users deserve to know when AI affects them
34	Can you provide audit logs of AI decisions? → Necessary for accountability and investigation
35	How do you handle requests to explain specific AI decisions? → Understand response time and thoroughness
36	Do you provide tools for monitoring AI behavior? → Dashboards, alerts, analytics

8 Performance and Reliability

37	What is your AI's accuracy rate? → Expect specific metrics — understand how accuracy varies across contexts
38	What are known failure modes? → Every AI has failure modes; transparent vendors acknowledge them
39	How do you handle errors and edge cases?

	→ Error handling matters as much as regular operation
40	What monitoring and alerting do you provide? → Proactive monitoring catches problems early
41	What are your SLAs and performance guarantees? → Understand commitments and remedies for failures

9	Business and Legal
42	What happens if I need to stop using your AI immediately? → Understand data retrieval and system suspension
43	Who is liable if your AI causes harm? → Understand liability allocation clearly
44	What insurance do you carry for AI-related claims? → Professional liability, errors and omissions, cyber insurance
45	Have customers sued you or made claims related to AI failures? → How they handled past issues indicates future handling
46	What happens if new regulations make your AI non-compliant? → Understand vendor commitment to ongoing compliance

<p>Thorough vendor due diligence prevents most AI reputational crises. Vendors who can't or won't answer these questions are ones you should probably avoid.</p>	<p>Reviewed by: _____</p> <p>Date: _____</p>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------