


5 Questions to Vet AI Advice Before You Trust It

Don't let hype drive your architecture. Use this framework to protect your systems—and your sanity.



1

Has this person ever deployed production software in a secure, enterprise environment?

Smart Advice:
Grounded in real-world experience and systems thinking

Red Flag:
Comes from someone who's only built demos or research prototypes




2

Does this advice account for compliance, security, or data privacy?

Smart Advice:
Mentions PCI, HIPAA, FedRAMP, or regulatory constraints


Red Flag:
Ignores risk, assumes no rules apply



Does the solution integrate into your current tech stack—or require starting over?

Smart Advice:
Works with your .NET systems and Azure infrastructure

Red Flag:
Requires replatforming, rewriting, or replacing your tools





4

Can this scale, be tested, logged, and maintained by your existing team?

Smart Advice:
Includes observability, versioning, and lifecycle planning

Red Flag: Built in notebooks, no CI/CD path, no rollback strategy



5

Will this advice pass an internal code review or security audit?

Smart Advice:
Aligns with internal controls and architectural standards

Red Flag: Uses shady plugins, vague endpoints, or unreviewed libraries

