CYBERAGI

PRINT IT · BRING IT TO THE DEMO

The on-prem AI security checklist

12 questions to ask before any vendor touches your findings.

Every AI security vendor will tell you they don't train on your data, that the model is "proprietary," and that pricing is "simple." None of those claims can be checked in a demo. These twelve questions can.

FOUR SECTIONS · TWELVE QUESTIONS

A

Where does my data actually go?

B

Whose model is it?

C

What does it actually replace?

D

What will it really cost?

SECTION A Where does my data actually go?

01 Trace one finding end to end. Which network hops sit outside my perimeter?

GOOD ANSWER

A diagram, drawn on the spot, that ends inside your network.

RED FLAG

"It's encrypted in transit" as the whole answer. Encryption is not location.

02 What in your architecture, not your contract, prevents you from training on my findings?

GOOD ANSWER

The model and the data never share a network with the vendor.

RED FLAG

A pointer to the Terms of Service. Terms change. Architecture doesn't.

03 Which telemetry, prompt logs, or traces leave my environment, and who at your company can read them?

GOOD ANSWER

A named list, or "none, and here is how you can verify that yourself."

RED FLAG

"Only anonymized metadata." Ask what anonymizes an internal hostname.

SECTION B Whose model is it?

04 Is the model yours, or a call to someone else's API with your interface on top?

GOOD ANSWER

Named model weights, running on hardware you can point to.

RED FLAG

"We use best-in-class foundation models." Ask whose, and where they run.

05 Does the model improve on my findings over time, and does that improvement stay on my hardware?

GOOD ANSWER

Fine-tuning happens on your box. The weights are yours to keep.

RED FLAG

"Learnings are shared across our customer base." That is your data leaving.

06 If your upstream model provider changes price, policy, or availability, what happens to my deployment?

GOOD ANSWER

Nothing. There is no upstream.

RED FLAG

"We'd migrate you." Ask how long that took the last time.

SECTION C What does it actually replace?

07 Which tools come off my invoice on day one, and which are you only sitting on top of?

☐

GOOD ANSWER

A list by name, including the tools you will still need to keep.

RED FLAG

"We integrate with everything." Integration adds a tool. It doesn't remove one.

08 Do the pentest, the SOC alert, and the threat model share one context, or three databases behind one login?

☐

GOOD ANSWER

A pentest finding appears in the threat model without anyone re-entering it.

RED FLAG

A separate dashboard per module.

09 Show me one unedited report the platform produced for a customer my size.

☐

GOOD ANSWER

A real artifact, redacted, that you could hand to an auditor today.

RED FLAG

A sample built for demos. Ask when it was generated.

SECTION D What will it really cost?

10 What is my total annual cost at 1 user and at 200 users, every module, every token?

☐

GOOD ANSWER

The same number twice.

RED FLAG

A pricing "framework." Metered pricing punishes the team for adopting the tool.

11 What does the next module or feature cost after I sign?

☐

GOOD ANSWER

Zero. It is included.

RED FLAG

"Available as an add-on." Budget for the add-ons you don't know about yet.

12 If I stop paying, what do I keep: my data, my reports, the model weights, the hardware?

☐

GOOD ANSWER

All four, or a written list of what leaves with them.

RED FLAG

Export "on request." Request it during the trial and time it.

HOW TO SCORE IT

Count the straight answers.

12	Twelve straight answers. Buy.
9–11	Negotiate the gaps into the contract before you sign, not after.
<9	The demo was the product.

> Excalibur, CyberAGI's on-prem AI security platform, was built to clear all twelve. If you want that proven live on your own findings, we will walk through every one.

BOOK THE DEMO

cal.com/cyberagi/discovery