Is our AI safe? The one-page board answer.

CompoundTrust

A starting framework from CompoundTrust, not a standard. Fill it in for one live AI agent, then the next.

AGENT OR AI USE _____ REPORT OWNER _____ DATE · NEXT REVIEW _____

THE ANSWER (TOP OF THE PYRAMID)

We are enabling AI safely. We can see our AI: _____. We control what it can expose: _____. We can show how fast we act: _____.

THREE BOARD LINES

1. AI uses listed with an owner: ___ of ___
2. Live agents with a signed permission set and a tested stop: ___ of ___
3. Last detection drill (date): what fired · who acknowledged · how long: _____

TOP THREE GAPS

1. _____ Owner _____ Date _____
2. _____ Owner _____ Date _____
3. _____ Owner _____ Date _____

PILLAR	EVIDENCE · WHAT TO SHOW	EVIDENCE HELD	LAST CHECKED	OWNER	GAP · DATE
Visibility	**Listed** Inventory row (Day 1): owner, purpose, what it reads, reaches and can do, whose identity it runs as				
	Told One-page rules (Day 15) and tools people declared last month				
Exposure	**Allowed** Permission set (Day 30): what it can read, where output goes, two signatures (owner and security)				
	Watched Last detection drill (Day 60): did the alert fire, who acknowledged it, how long it took				
Defensibility	**Timed** Six clocks (Day 45): slowest result for revoke; has stop been tested				
	Frameworks Each applicable framework mapped to the evidence above: EU AI Act (where it applies), NIST AI RMF, privacy rules. AI vendors: current SOC 2 Type II report?				

READ THIS A gap is allowed. An undated gap is not. Every gap gets an owner and a date.

NOT THE HEADLINE

Counts of policies, training completion, tools bought and blocked-data volumes may matter, but they show activity, not control. Use them as supporting detail.

SPEAK RISK, NOT MODELS

Lead with material risk, regulatory exposure and business enablement. Keep prompt injection and model detail in the appendix.

WHAT I WOULD DO THIS WEEK

Pick one live agent. Fill in the six rows. Put a date on every gap. Take this page to the next leadership meeting.

Builds on the Season 2 kit pages: AI inventory (Day 1), acceptable use (Day 15), agent permissions (Day 30), six clocks (Day 45), misuse detection (Day 60). Framework applicability is a judgement for your legal and compliance teams.