THE AMPLIFI AI SECURITY FRAMEWORK

We teach AI *responsibly.* Here's how.

Security isn't a compliance checkbox — it's a teaching posture. Every programme we run, every engagement we deliver, and every tool we recommend is filtered through five principles. This is the full framework.

AMPLIFI AI · A GLOBAL AI CONSULTANCY

Bridgetown, Barbados · Serving 32 markets · amplifiai.co
Founded by Janelle Germain, Founder & CEO

THE FRAME

AI is too powerful

to use without principles.

Most AI security conversations stop at the tool level. “Is ChatGPT secure?” “Is Claude enterprise-grade?” Those questions matter — but they’re the wrong place to start. The more useful question is: what are you putting into the tool, and what are you doing with what comes out?

Every Amplifi AI engagement — whether you’re in the Masterclass, running a Corporate cohort, or working with us on Consulting or Studio — is shaped by a shared framework. Five principles that apply regardless of which model you use, which tier you pay for, or which jurisdiction you operate in. These principles are how we teach AI security, how we govern our own practice, and how we expect clients to govern theirs.

What follows is the full framework. Read it once to understand the shape. Come back when you need to teach it to your team, draft a policy, or answer a board question about AI risk.

THE FIVE PRINCIPLES

Five rules that scale from

solo operator to enterprise team.

01

Never paste customer data into free AI tools.

WHY THIS MATTERS

Free tiers of AI tools — ChatGPT free, Claude free, Gemini free — train on your inputs by default. Customer names, email addresses, transaction records, support tickets, medical details, financial data: all of it can be ingested into the model's training corpus if you paste it into a free-tier chat. Once it's in, you can't get it out.

WHAT THIS LOOKS LIKE IN PRACTICE

Before pasting anything into an AI tool, ask: could this data identify a customer, a patient, an employee, or a minor? If yes, use a paid tier with data retention controls turned off. If that option doesn't exist, redact identifying fields or use synthetic data instead.

RED FLAGS

- Team members uploading spreadsheets of customer records to free AI tools for 'summarisation'
- Pasting transaction data, support emails, or medical records without checking the tool's data policy
- Using free AI tools to 'clean up' lists that include real names, emails, or identifiers
- Assuming that deleting the chat deletes the data — it doesn't

EXAMPLE

A marketing team exports 2,000 customer emails into a free AI tool to generate 'personalised' outreach subject lines. Those 2,000 emails are now permanent training data. This has happened to brands that paid dearly for it.

02

Never share confidential business information with unverified AI.

WHY THIS MATTERS

Unreleased product plans, financial projections, strategic roadmaps, M&A discussions, employee compensation, legal exposure: these are the exact categories of information that competitors, regulators, and journalists would pay for. Pasting any of it into an unverified AI tool is the same as emailing it to an unknown address.

WHAT THIS LOOKS LIKE IN PRACTICE

Before using any new AI tool for strategic work, verify three things: (1) the company behind the tool, (2) the data retention policy in writing, and (3) whether the tool is on your organisation's approved list. If you can't verify all three, default to not pasting anything confidential.

RED FLAGS

- Browser extensions that 'summarise with AI' without disclosing where the data is sent
- 'AI assistants' built by unknown startups with no published security policy
- Using personal accounts on work-adjacent AI tools because the enterprise account is slower
- Assuming that a tool is safe because 'everyone uses it'

EXAMPLE

An executive pastes a draft acquisition memo into a browser-extension AI summariser to get the gist. The extension routes the request through a third-party API with no data controls. The memo is now in someone else's logs.

03

Understand the difference between free and paid AI tiers.

WHY THIS MATTERS

Free and paid tiers of the same AI tool are not equivalent products. They differ in data handling, retention, training use, and enterprise controls. Using the free tier for work-related tasks is a common shortcut that creates uncommon risk.

WHAT THIS LOOKS LIKE IN PRACTICE

Know, for every AI tool your organisation uses: (1) what the paid tier offers that the free tier doesn't, (2) whether your team is on the right tier for the work they're doing, and (3) what data the free tier retains for training. If the delta matters, budget for the paid tier — it's almost always cheaper than the alternative.

RED FLAGS

- Using a personal free-tier account for work because 'IT is slow to approve the paid version'
- No clarity about which team members are on which tier
- Assuming the paid tier is automatically private — most require explicit admin settings to disable training
- Free tier nested inside a bigger paid account (e.g. team member on free tier despite org paying for Teams)

EXAMPLE

A senior manager uses the free version of an AI tool to draft an internal memo because 'it's just for internal use.' The free tier trains on the memo. Six months later, a competitor's AI assistant knows the memo's key phrases verbatim.

04

Know your jurisdiction's data protection laws.

WHY THIS MATTERS

GDPR, CCPA, HIPAA, PIPEDA, PDPA, LGPD — every jurisdiction has rules about where personal data can go, who can process it, and what happens if it leaks. AI tools don't know your jurisdiction. You do. The legal exposure when data crosses borders or gets processed by a third-party AI is on your organisation, not the tool vendor.

WHAT THIS LOOKS LIKE IN PRACTICE

Before deploying AI at scale in any function that touches personal data, answer three questions in writing: (1) which jurisdictions do our customers live in, (2) what does each jurisdiction require for AI processing of personal data, and (3) where does our AI vendor's infrastructure process and store data. If the answers don't align, you have a compliance problem — not an AI problem.

RED FLAGS

- Deploying AI customer-service tools without checking if your AI vendor is in your customers' approved jurisdictions
- EU customer data being processed by AI models whose servers are outside the EEA
- No Data Processing Agreement with your AI vendor
- Marketing automations using AI to process customer data without explicit consent disclosure

EXAMPLE

A retail brand deploys an AI chatbot for EU customer support. The bot sends every conversation to a US-hosted model. No DPA exists. The brand is now in breach of GDPR Article 44 and doesn't know it.

05

Always review AI-generated output before publishing.

WHY THIS MATTERS

AI models hallucinate — they confidently produce information that is false, defamatory, biased, or legally exposed. Publishing AI output without human review puts the organisation on the hook for everything the model says. 'The AI wrote it' is not a legal defence.

WHAT THIS LOOKS LIKE IN PRACTICE

Every piece of AI-generated content that leaves the organisation — public posts, client emails, reports, code, marketing copy, customer-facing documentation — passes through a named human reviewer before publication. The reviewer's job is not to approve; it's to catch hallucinations, verify claims, and take responsibility for the output.

RED FLAGS

- AI-generated content auto-publishing to social or public channels
- No named owner for AI-generated outputs in any given workflow
- Reviewers skimming AI output rather than fact-checking it
- Using AI to generate citations, quotes, or statistics without verifying sources

EXAMPLE

A marketing team publishes an AI-drafted blog post with a fabricated statistic attributed to a real research firm. The research firm issues a legal demand. The brand pulls the post, publishes a correction, and absorbs reputational damage that a 10-minute fact-check would have prevented.

IN PRACTICE

How we apply this framework

across Amplifi AI.

The framework isn't theoretical. It shapes the day-to-day across every Amplifi AI arm.

AMPLIFI ACADEMY

Every programme starts with the framework. Masterclass introduces the five principles; Pro teaches you how to write the governance policies, security tiers, and incident-response protocols that turn principles into operating systems.

AMPLIFI CORPORATE

Corporate engagements include a governance module mapped to your industry's regulatory context — financial services, healthcare, retail, SaaS, or regulated jurisdictions like the EU.

AMPLIFI CONSULTING

Consulting engagements typically begin with a governance audit — we map your current AI usage against the five principles and identify the highest-risk gaps first.

AMPLIFI STUDIO

Every Studio deliverable — website, app, campaign, brand system — is engineered with AI guardrails in place. Nothing we build leaks training data, embeds unverified AI, or bypasses client-approved tiers.

AMPLIFI TALKS

Governance is the most-requested topic for keynotes and corporate workshops. If you've booked a Talk on 'AI adoption,' expect the framework to show up.

GO DEEPER

Train your team on the framework.

The five principles are the foundation. The full curriculum — governance frameworks, security tiers, policy templates, incident-response protocols — is taught in Amplifi Academy Pro and delivered in-depth to Corporate cohorts.

Amplifi Academy Pro · 4-week strategist cohort · amplifiai.co/pro

Amplifi Corporate · Customised team training · amplifiai.co/corporate

Questions? · hello@amplifiai.co