

FOR ALARM COMPANIES · SEPTEMBER 2026

Safe AI with Customer Data

A decision guide: which AI, for which data, and the five questions to ask before it touches a customer record.

NAOMI WITHERS, EMBA, CHPC · GROWTH CONSULTANT SERVICES

WHERE TO START

The data decides the tier. Not the brand of AI.

Every major AI lab now sells a plan that keeps your data out of model training. Most exposure comes from putting the wrong data into the wrong plan. So we start with your data, not the tools.

WHY OUR INDUSTRY IS DIFFERENT

Your data is not ordinary customer data

RED • NEVER

Keys to the building

- Passcodes and verbal passwords
- Alarm, gate and lock codes
- Zone lists and layouts
- Open and close schedules

Stays in your system of record. Never pasted into any AI tool.

AMBER • BUSINESS PLAN ONLY

Who and where

- Names, addresses, phones
- Emergency contacts
- Billing and account details
- Interior video (platform AI only)

Only in a business plan with a data agreement, or inside your software's own AI.

GREEN • OPEN

How you work

- Procedures and install standards
- Marketing and proposals
- Training material
- Public information

Any business-plan AI tool is fine.

A leak of ordinary customer relationship management (CRM) data leads to phishing. A leak of Red data can lead to a burglary.

WHAT HAS ALREADY HAPPENED

The door was people and connected apps, not the AI model

Two national alarm companies

breached in 2026. In both, an attacker phoned an employee, took over their login, and pulled customer records from the CRM.

700+ organizations

exposed in 2025 when attackers stole the access keys of an AI chatbot connected to Salesforce.

Three major chatbots

had users' shared chat links show up in Google between 2025 and 2026, including personal details and passwords.

AI adds two new doors: **what your people paste in**, and **what connected AI apps can reach**.

THE FOUR LEVELS OF ACCESS

Paid is not the same as business

Level	Examples	Trains on your data?	Staff may read chats?	Customer data?
Free personal	Personal accounts at any lab	Yes, by default	Yes	No
Paid personal	Plus, Pro, Max, Google AI Pro, SuperGrok	Yes, by default	Yes	No
Business / Team	ChatGPT Business, Claude Team, Workspace Gemini, Microsoft 365 Copilot	No, by default	Abuse checks only	Amber data, with a policy
Enterprise / API	Enterprise plans, API, and the AI inside your software	No	Abuse checks only	Yes, with governance

The rule: always a business plan, on a work account. Commercial terms and a data processing agreement are what protect you, not the price.

The same pattern at every lab

Lab	Personal plans (free and paid)	Business plan	Enterprise / API
OpenAI	Free, Go, Plus, Pro: trains by default; opt out	Business: \$20–25/user, 2-seat min, no training	Enterprise, API: no training; Business Associate Agreement available
Anthropic	Free, Pro, Max: trains by default; opt out	Team: \$20–25/user, 2-seat min, no training	Enterprise, API: no training; Business Associate Agreement available
Google	Gemini app: trains by default; reviewers may keep chats 3 years	Gemini in Workspace: no training, no human review	Vertex, paid API: no training (free API trains)
xAI	Grok, SuperGrok: trains by default; opt out	Grok Business: no training; data agreement terms unpublished	Enterprise, API: no training
Microsoft	Personal Copilot: trains by default; opt out	Copilot Chat, work account: no training, \$0 extra	Microsoft 365 Copilot, \$21–30/user: no training

A Business Associate Agreement is the contract U.S. health-privacy law (HIPAA) requires before a vendor may handle health information. Terms and list prices as of September 22, 2026; check vendor pages before you buy.

THE EASIEST SAFE FIRST STEP

It may already be paid for

ON MICROSOFT 365?

Copilot Chat, signed in with a work account

Enterprise data protection, no training, no extra cost on eligible plans.

ON GOOGLE WORKSPACE?

Gemini inside Workspace

Covered by your Workspace terms and admin controls. No training, no human review.

EITHER WAY

Close the side doors

Turn off public share links in the admin console. Treat any share link as public. Review browser extensions.

One caution: this kind of AI sees everything the user can see. Tidy up file and folder permissions first.

A COST-BENEFIT QUESTION

Should an AI company learn from your data?

WHAT YOU GET

- Sometimes free usage. One lab offers qualifying API customers up to 1M free tokens a day for sharing their traffic
- Better models over time, for everyone
- Nothing to set up. It is the default

WHAT IT COSTS YOU

- It cannot be taken back. Opting out later does not remove data already used for training
- Long retention: up to 5 years on one lab's personal plans
- Legal and customer-trust exposure. Consumer chat logs have already been pulled into lawsuits

For customer data, the answer is always no. For generic work, it is your call. Either way, look for the setting in every tool you use.

ONE LINK THAT STAYS CURRENT

Safety and privacy are two different questions

SAFETY OF THE COMPANY

Future of Life Institute AI Safety Index

An independent nonprofit panel grades the major labs twice a year. The link always opens the latest edition.

futureoflife.org/index

PRIVACY OF YOUR DATA

Ask for the evidence

No independent privacy rating stays current. Ask for the vendor's trust center, its System and Organization Controls (SOC 2) Type II audit report, its ISO/IEC 42001 AI certification, and “no training” in the contract.


Grades are relative. In July 2026 no lab scored above C+. Compare labs, and watch the direction.

The two can disagree. One lab graded F on safety ranked first in a 2025 privacy ranking.

Check who and when. Trust independent ratings over vendor claims. Treat anything older than a year as stale.

OPEN-SOURCE (OPEN-WEIGHT) MODELS

Lowest safety grades, most private option



Future of Life Institute AI Safety Index, Summer 2026 (July 2026)

THE PARADOX

Open-weight labs grade lowest on safety. But a model you run on your own computer never sends customer data anywhere.

THE TRADE-OFF

The best open models trail the leaders by about four months. Models small enough for one workstation trail further, and you own the upkeep.

Six rules for Personally Identifiable Information (PII)

1 Minimize

Send only the fields the task needs.
Red data never goes.

2 Keep it in the system

If the data lives in your customer system or platform, use that vendor's governed AI.

3 Swap names for tokens

Replace names with codes like CUST_0142; map them back yourself.

4 Test with fake data

Use synthetic customer records for pilots, demos and training.

5 Process raw PII locally

A local model strips PII before a frontier model sees the text.

6 Tell customers openly

If your AI use changes, update your privacy notice openly.

THE SOFTWARE YOU ALREADY RUN

Check the defaults in your software

CRM and sales

Mature, and often already switched on

Back office and field service

Arriving now as add-ons and credits

Monitoring and platforms


In production for video verification and false-alarm triage

Examples	What the vendor's own help pages say (as of September 2026)
HubSpot	Trains its own AI on customer data by default. Only a Super Admin can opt out.
ServiceTitan	Customers are opted in to data sharing for product and AI development by default.
Salesforce	Zero retention with its AI providers, but data masking is off when agents are used.

Understand your defaults and your settings. Ask every vendor: Is it on by default? Does it train on our data? Which AI provider sits underneath, and what data does it keep?

THE CUSTOMIZATION LADDER

From \$20 a month to \$100 million



Log scale: each gridline is ten times the one before. Sources: vendor pricing pages; Compute Comparison (2026); Epoch AI; Stanford AI Index. Ranges are approximate.

Cheap compute, costly data

THE COMPUTING: CHEAPER THAN YOU THINK

- About \$2–\$4 per training run on a small open model
- About \$45–\$120 on a large one
- Far below the \$40M+ to train a model from scratch

THE DATA AND UPKEEP: HARDER THAN IT LOOKS

- \$3K–\$26K to label examples by hand
- An engineer to test it, and redo it when the base model is retired
- Models can memorize and repeat training data. Never train on raw customer PII

Worth it only for a narrow, high-volume, repetitive task, with PII removed first. To teach AI about your business, retrieval over your documents is cheaper and easier to update.

BUILD VERSUS BUY

Buy or wait before you build

BUY

Under \$10K a year

Business-plan seats for a 25-person team, plus Copilot for power users. Working in weeks.

WAIT

Let your vendor add it

Gartner expects 40% of enterprise apps to include AI agents by the end of 2026. Then check the defaults.

BUILD

Only with a reason

\$15K–\$120K to build, plus upkeep. Three-year cost often runs two to three times the build.

In one 2025 MIT study, purchased AI tools succeeded about twice as often as internal builds. Treat it as directional, not precise.

Five questions to ask first

1	What data will it see?	Red, Amber or Green. Red never goes into a general tool.
2	Whose terms govern it?	Consumer or business? A data agreement? “No training” as the default, in writing?
3	Where does it go, and for how long?	Retention, zero-retention options, which AI provider sits underneath, and where.
4	Who else can reach it?	Share links, connectors, permissions. Multi-factor login everywhere.
5	Who is accountable?	A named owner, a one-page policy, a person in the loop, and an incident plan.

WHEN GOOD HABITS ARE NOT ENOUGH

Any one of these means formal governance

- **Red data is involved.** Codes, passwords, schedules, zone lists.
- **The AI can take actions.** Writing records, sending messages, changing billing.
- **It influences a decision about a customer.** Credit, financing, cancellations. California's automated-decision rules start January 1, 2027.
- **Video, biometrics or face recognition.** State biometric laws apply, such as Illinois'.
- **You are building or connecting something custom.** Retrieval, integrations, fine-tuning, local models.

FRAMEWORKS TO LEAN ON

National Institute of Standards and Technology (NIST) AI Risk Management Framework

ISO/IEC 42001, the international AI management standard, built to scale down to smaller firms

Cybersecurity and Infrastructure Security Agency (CISA) AI Data Security guidance (May 2025)

BEFORE THE END OF THE MONTH

Three things to do

1

Move to business plans

Everyone on a business plan with a work account. Training off. Public share links off.

2

Write your one-page rule

Red, Amber, Green, adapted to your data. Name an owner.

3

Open your settings

In your CRM and platforms, check the AI defaults and ask the three vendor questions.

Bookmark: futureoflife.org/index for safety grades, and your vendors' trust centers for privacy.

NAOMI WITHERS, EMBA, CHPC · GROWTHCONSULTANTSERVICES.COM

GROWTH CONSULTANT SERVICES

More free tools for alarm companies

growthconsultantservices.com/security

The Six Numbers calculator

Price a vendor's promise against your own loaded costs.

[/security/six-numbers](https://growthconsultantservices.com/security/six-numbers)

Alarm Company Valuation Estimator

What a buyer would pay for your recurring monthly revenue, and what your contracts are costing you.

[/security/alarm-valuation](https://growthconsultantservices.com/security/alarm-valuation)

AI Transformation Diagnostic

Where your organization stands across five dimensions of AI readiness.

[/ai-transformation](https://growthconsultantservices.com/ai-transformation)

Naomi Withers, EMBA, CHPC, founded Growth Consultant Services after more than a decade of finance and operations leadership inside a top SDM 100-ranked security company. She advises founders and leadership teams on AI strategy and organizational readiness.

hello@growthconsultantservices.com

Updated September 2026. Educational, not legal advice. Vendor terms and prices change; confirm current terms with each vendor.