



AI Safety for Ski Areas

What to Allow, What to Block, and How to Write a Policy

WHAT WE'LL COVER · 30 minutes

- 01 Where the real risks are** *and where they aren't*
- 02 The 3-Question Test** *before any AI tool touches your data*
- 03 Specific rules for your data** *customer, employee data*
- 04 DEMO: Big 3 settings comparison** *ChatGPT · Claude · Gemini*
- 05 Vendor questions + starter policy** *what to ask, what to write down*

drinkcaffeine.com

DISCLAIMER



We Are Not Your Attorneys

INFORMATION PROVIDED IS OUR PERSPECTIVE
FROM WORKING WITH DOZENS OF COMPANIES

The logo for 'caffeine' is displayed in a green, lowercase, sans-serif font. Above the letter 'i' in 'caffeine' is a small, stylized graphic of a coffee bean or a cluster of three dots.

WHY THIS MATTERS NOW

AI now has business AND legal risks.



February 2026 — U.S. v. Heppner

Federal court ruled AI conversations on consumer tiers carry no attorney-client privilege and no work-product protection. Discoverable in litigation.



Anthropic flipped Claude's default in Sept 2025

Consumer Claude (Free/Pro/Max) now defaults to training on your data, with 5-year retention. The "privacy-first" provider quietly changed sides.



Your insurance carrier is asking

AI governance is now appearing on cyber liability renewals and motorcoach operator audits.

caffeine

WHERE THE REAL RISKS ARE

This is what actually hurts you

- Employees pasting customer information into consumer ChatGPT
- Employee records used as test data in a free AI tool
- Pricing & staffing leaking into training data for competitors
- Vendor AI features your team enabled without IT or legal review
- AI-generated content (documents, emails) sent without verification

THE FRAMEWORK

The 3-Question Test- Internal



1

Where does our data go?

And how long does it stay there?

- Is it used to train the model?
- How long is it retained?
- Who else can see it (humans, contractors)?
- Can we delete it on demand?

2

What's the worst case if it leaks?

Pretend it's posted publicly tomorrow.

- Is it PII (customer names, payment info)?
- Is it regulated (driver HOS, DOT records)?
- Is it competitive (pricing, contracts)?
- Can it harm a person, not just the company?

3

Who's accountable when it breaks?

AI vendors love to disclaim everything.

- Who owns the output legally?
- Are there indemnification clauses?
- Is there a DPA (Data Processing Addendum)?
- Does our insurance cover AI-caused errors?

The 3-Question Test- Vendors



1

Where does our data go?

Q1 on the vendor questionnaire.

- Which model provider(s) do you use? (OpenAI, Anthropic, Google, others)
- Are you on their consumer, business, or enterprise API tier?
- Where is data processed and stored geographically?
- Will our data be used to train the underlying model? Get this in writing.

2

What if it leaks?

Push for real answers, not badges.

- Do you have SOC 2 Type 2? Provide the report, not just the badge.
- What's your data breach notification timeline?
- Have you had a breach in the past 24 months? What happened?
- What happens to our data when we terminate the contract?

3

Who's accountable?

Contracts and insurance matter.

- Does the contract have a Data Processing Addendum (DPA)?
- What's your AI-specific indemnification language?
- Who owns the outputs your AI generates from our data?
- What's your cyber liability coverage limit?

LIVE DEMO



The Big 3 Settings Walkthrough

caffeine

THE BIG 3 • REFERENCE TABLE

ChatGPT vs. Claude vs. Gemini — the settings that matter

Configuration	ChatGPT (OpenAI)	Claude (Anthropic)	Gemini (Google)
Consumer default: trains on data?	YES (Free/Plus/Pro)	YES (Free/Pro/Max) since Sept 2025	YES (consumer) plus 3-yr human review
Where to turn off training	Settings → Data Controls → "Improve the model for everyone"	Settings → Privacy → "Help improve Claude"	Settings→ Activity→ Keep activity
Retention if you DO opt in	Indefinite (up to 30-day abuse logs after deletion)	Up to 5 years (60x longer than default)	3-yr human review retention
Business / Team tier name	ChatGPT Team / Enterprise	Claude Team / Enterprise (Claude for Work)	Google Workspace with Gemini
Business default: trained on?	NO — contractually	NO — contractually	NO — contractually
SAML SSO / admin console	Yes (Team/Ent)	Yes (Team/Ent)	Yes (Workspace)
HIPAA / BAA available?	Yes (specific tiers)	Yes (commercial)	Yes (Workspace + BAA)
Best for motorcoach operators if...	You want broadest 3rd-party tool integration	You want strongest default privacy on the API	You already run Google Workspace

All three providers offer contractually safe business tiers. The dangerous version is the consumer tier your team is using on their personal account.





DATA-SPECIFIC GUIDANCE • 1 OF 3

Customer Data



ALWAYS BLOCK

- Full passenger manifests
- Payment / card data
- Driver license numbers
- Email lists with names attached
- Anything you wouldn't put in a Google review

OK WITH GUARDRAILS

- Anonymized trip patterns
- Aggregated booking trends (no names)
- Drafting templates from de-identified examples
- Sample/dummy data for testing

GENERALLY SAFE

- Industry research questions
- Drafting marketing copy
- Summarizing public documents
- General "how do I..." questions

Rule of thumb: *If you wouldn't email it to a competitor, don't paste it into a consumer AI tool.*

caffeine



DATA-SPECIFIC GUIDANCE • 2 OF 3

Employee & Guest Records



This is the category with the most regulatory exposure

Guest injury records, employee medical files, and drug & alcohol results are protected under HIPAA, state ski-safety statutes, and state privacy laws. AI tools that retain or train on this data create exposure on all fronts.

ALWAYS BLOCK

- Patrol incident reports with guest names
- Drug & alcohol test results
- Employee medical / fitness-for-duty records
- HR & disciplinary files
- Guest injury / accident details with names

OK WITH GUARDRAILS

- Aggregated on-mountain incident patterns
- De-identified injury summaries for safety training
- Staff scheduling logic (no names)
- Policy & procedure templates

GENERALLY SAFE

- State ski-safety statute summaries
- Industry best-practice questions
- Drafting an employee or guest-services handbook
- General compliance “how do I...” questions

Use the Team or Enterprise tier of your AI tool for anything in the middle column. Never the consumer tier.

caffeine



Pricing & Quoting Information



Why this matters: pricing data is competitive intelligence

Your dynamic lift pricing, season pass tiers, group discounts, and corporate account rates are the most strategically valuable data you have. They are also the most commonly pasted into AI tools for “quick help.”

ALWAYS BLOCK

- Named corporate / group account rates
- Live or pending RFP & event bids
- Pricing strategy vs. specific competing resorts
- Margin / cost structures per lift ticket or pass
- Negotiated B2B contracts & MFN clauses

OK WITH GUARDRAILS

- Published lift ticket & season pass ranges
- Generic dynamic-pricing models (no account names)
- Group & lesson quote templates
- Standard add-on / rental fee calculations

GENERALLY SAFE

- Market research on competitor resort pricing
- General revenue-management methodology questions
- Industry trend analysis (visits, snowfall, demand)
- Drafting a public rate sheet or window price list



CONNECTORS SECURITY



Your Data, Their Servers - What Gets Exposed

NEVER ENABLE

- Connectors on consumer (Free/Pro/Max) accounts
- Connecting your CRM, email, or HR system to a personal Claude account
- Any connector with access to regulated data (passenger PII, financial records)
- Auth tokens with write/delete permissions — no grants without IT review

ENABLE WITH LIMITS

- Google Drive, Slack, or Calendar — on Team/Enterprise tier only
- Read-only OAuth scopes — no write/delete permissions granted
- IT-approved connector list reviewed quarterly; connector added by IT admin only
- Signed DPA in place with Anthropic confirming no training on connected data

SAFE TO ENABLE

- Public web search and research tools
- Industry databases with no internal company data
- Connectors on Team/Enterprise with confirmed data residency and no training opt-in



Rule of thumb:

A connector is a permanent open window into your data. Treat it like a vendor contract — require IT sign-off, read-only access, and a signed DPA.





AGENT SECURITY



AI on Your Desktop - What It Can See and Do

NEVER ALLOW

- Running agent on a machine with access to payroll, HR, or legal systems
- Granting agent the ability to send emails or submit forms autonomously
- Using agent on a consumer account — your screen activity feeds the LLM's training data
- Installing on shared/admin machines without a prior security review

USE WITH CONTROLS

- Team/Enterprise tier only — confirm “Help improve XYZ” is toggled OFF before install
- User must approve every action before agent executes it — no fully autonomous mode
- Deploy only on isolated, work-dedicated machines — not personal laptops with personal data
- IT logs all agent sessions; review access scope before rollout

GENERALLY SAFE

- Public web search and research tools
- Industry databases with no internal company data
- Connectors on Team/Enterprise with confirmed data residency and no training opt-in

APPLYING GUARDRAILS



You've identified “Guarded data.” Now what?

1

Strip the identifiers

Replace names, pass numbers, and IDs with generic labels before typing anything into an AI tool. “Guest A”, “Employee 3”, “Pass Holder X” — not real names.

2

Describe the pattern, don't paste the record

Tell the AI what you need help with in plain English. Don't paste a raw incident report, reservation record, or guest history. If you can't describe it without pasting it, it's probably a Block.

3

Confirm you're in the right tool

Check the account name in the corners of your screen. If it's a personal email address, stop. You need the company account.

4

Review before you use the output

Don't send AI-generated summaries of guest or employee data to anyone — internally or externally — without reading it yourself first.

WORKED EXAMPLE: Guest incident pattern

DON'T

“Here is the incident report for Jane Doe, pass #SK-20291. She was injured on 10/4 on run 7. Her medical release is attached. Can you write a liability summary?”

DO

“I have a guest (call them Guest A) who was injured on an intermediate run. Can you help me draft a neutral incident summary template for our legal team?”

Key difference: the DO version gets the same useful output with zero identifiable guest data in the prompt.

The honest test:

If you can't describe what you need without pasting the actual record — stop. It's a Block, not a Guarded.

caffeine

PROTECT YOURSELF



Your 1-page AI usage policy

1. Purpose & Scope

Applies to all employees, contractors, and vendors who use AI tools on company business or with company data.

3. Data Classification

Block: PII, driver records, pricing, contracts. Guarded: anonymized data. Open: public / industry info.

5. Vendor AI Requirements

Any vendor using AI on our data must sign a DPA, name their model provider, and confirm we are not in training pipelines.

7. Incident Reporting

Any suspected AI-related data exposure reported to [role] within 24 hours. No retaliation for self-reporting.

2. Approved Tools

List specific business-tier AI tools by name. Consumer / personal accounts are not approved for company work, regardless of who pays.

4. Required Settings

Training opt-out must be confirmed on every approved tool. IT verifies quarterly. Personal accounts are not used for company data, period.

6. Output Verification

AI-generated content sent externally (quotes, contracts, customer comms) requires human review before sending.

8. Review Cycle

Policy reviewed semi-annually or when a provider changes its data policy (e.g., Anthropic's Sept 2025 change).



GET STARTED

30 Day Action Plan



WEEK 1

Audit

- Survey your team: which AI tools are they actually using?
- Check whether any are consumer (Free/Plus/Pro/Max) accounts touching company data
- Inventory vendors who added AI features in the past 18 months

WEEK 2

Configure

- On approved tools: verify training opt-out toggles are OFF
- Migrate critical workflows from consumer to Team/Business tiers
- Document settings in a shared IT checklist

WEEKS 3–4

Codify

- Adapt the starter policy to your operation
- Send vendor questionnaire to top 5 vendors
- Train staff: 15-minute “lunch & learn” on the 3-Question Test





Questions?

Three things to take with you:

- 1 Move company AI work to Team/Business tiers — consumer plans are a liability.
- 2 Apply the 3-Question Test to every tool, internal or vendor.
- 3 Adopt a written policy this quarter — your insurance carrier will ask.

drinkcaffeine.com

REQUEST THE DECK
+ AI SAFETY GUIDE



drinkcaffeine.com/intelligence

