VERMONT AI SYSTEMS · PRIVATE AI FOR BUSINESSES THAT CANNOT COMPROMISE ON DATA PRIVACY

The Private AI Buyer’s Guide

For Vermont businesses that need AI that works without handing your data to OpenAI, Anthropic, or Google.

A practical guide to evaluating private AI solutions and understanding what every vendor isn’t telling you.

FREE GUIDE · 8 PAGES · NO EMAIL REQUIRED TO READ

What’s inside:

SECTION 1 — THE PROBLEM

Public AI was built for consumers. Your business data wasn’t the product.

Every time an employee pastes client data into ChatGPT, submits a document to Claude, or uses Gemini on company information, that data potentially becomes part of a training corpus. Even vendors with opt-out policies have buried limitations, ambiguous data retention clauses, and infrastructure dependencies on the same public cloud providers that built the models.

The legal and competitive risk is real. The business decisions made around AI data security in 2024 and 2025 will shape compliance exposure for years. Most businesses are not having the right conversations with their IT teams, legal counsel, or AI vendors.

WHAT MOST BUSINESSES ARE GETTING WRONG

- "We have an NDA with our AI vendor" — NDAs don’t govern what happens inside a third-party model’s training pipeline.
- "We’re using the enterprise plan, so our data is protected" — Enterprise plans typically reduce retention windows, not eliminate them.
- "Our employees know not to share sensitive data" — You’re one pasted document away from an exposure, and you won’t know it happened.
- "We don’t have anything worth stealing" — Your business data, client relationships, and internal processes are worth everything to your competitors.

THE STAKES ARE REAL

HIPAA penalties for unauthorized PHI disclosure reach $1.9M per violation category per year. Attorney-client privilege can be waived when communications pass through third-party AI systems. Trade secrets lose protection the moment they enter a system outside your control.

The question is not whether your data is valuable. The question is whether your current AI tools are designed to protect it. Most are not.

WHAT MOST AI VENDORS WON'T TELL YOU

- Where training data actually comes from and whether your inputs could be included
- How long your data is retained, even under an enterprise contract
- What infrastructure the model actually runs on (hint: it's often a public cloud)
- Whether your data is used to improve the model for other customers
- What happens to your data if the vendor is acquired or changes terms

SECTION 2 — WHAT PRIVATE AI ACTUALLY MEANS

Private AI is not a feature. It’s an architecture.

There is no legal standard for "private AI." Any vendor can claim their product is private. Understanding what that means in practice requires examining three specific things: where the model runs, who has access to your data, and what happens to your data over time.

WHERE THE MODEL LIVES

A truly private AI runs on infrastructure you control — your own on-premise servers, a private cloud environment you own, or a dedicated VPC that is never co-tenanted with other organizations. When a vendor says "your data never leaves your environment," ask them to define the environment. If it includes a third-party API call, it is not private.

WHO HAS ACCESS TO YOUR DATA

During the build phase, the engineering team needs some access to your data to train and test the model. After deployment, that access should be entirely yours. Ask vendors to specify exactly which personnel have data access, under what conditions, and how access is logged and revoked. If they cannot produce an answer, that is a significant red flag.

WHAT "TRAINED EXCLUSIVELY ON YOUR DATA" ACTUALLY MEANS

This phrase is frequently used as marketing. In practice, it should mean: the model is fine-tuned on your documents and knowledge base only. Your data does not improve any public model. Your data does not leave your environment. Your data is never transmitted to a third-party API. Any vendor using this phrase should be able to explain exactly how each of those statements is technically true.

IP OWNERSHIP — THE QUESTION MOST BUSINESSES SKIP

When we build a private AI for you, you own it. The trained model weights, the fine-tuning dataset, the deployment configuration — all of it transfers to you at project close. We retain no copies. If you cancel the retainer, you keep running your own AI.

That means no vendor lock-in, no licensing dependency, no "but the model lives in our cloud"

clauses. If the vendor is not willing to put that in writing, the AI is not truly yours.

THE VERMONT AI SYSTEMS PRIVACY COMMITMENT

- Your data never touches OpenAI, Anthropic, or any public LLM API in the data pipeline
- No employee of Vermont AI Systems has standing access to your model or data post-deployment
- Access during the build phase is scoped, logged, and revoked at project close
- You own the model, the weights, the fine-tuning data, and the deployment configuration
- We design systems that are safe even if the network perimeter is breached

SECTION 3 — HOW WE WORK

Five stages. No surprises.

We've built private AI deployments for law firms, manufacturers, and healthcare practices. Every engagement starts the same way — with us listening before we write a line of code.

01

Discovery Call

30 min · Free

A conversation to understand your business, your data, and what you want AI to actually do. We'll tell you plainly whether private AI makes sense for your situation and what it would cost. No pitch, no pressure. If we're not the right fit, we'll say so.

02

Data Assessment

1–2 weeks · Included

We audit the data you have — documents, databases, emails, SOPs, CRM exports — and produce a written assessment of what we can train on and what gaps need filling. This gives you a complete picture before any contract is signed.

03

Private Model Build

3–8 weeks

We build and fine-tune your AI on your data, in a private environment we control. Your data never touches a public LLM. We run the model by you with real prompts before you commit to the full deployment.

04

Deployment & Staff Training

1–2 weeks · Included

We deploy to your environment — on-premise server, private cloud, or secure VPC — and run hands-on training with every employee who will use the AI. We don't hand you a user manual and disappear.

05

Ongoing Tuning

Monthly retainer · Optional

Your business changes. Your AI should too. Retainer clients get quarterly model refreshes, usage reviews, and direct access to our team. We track output metrics so you can see exactly what you're getting.

WHAT YOU OWN AT THE END

At project close, you receive: the trained model weights, your original fine-tuning dataset, the full deployment configuration, and a technical handoff document. Your AI runs on your infrastructure. We hand you the keys and you keep them.

There is no licensing fee, no per-user charge, no per-query billing. The model is yours. You run it. You own it.

SECTION 4 — SERVICE TIERS

Predictable scope. Honest pricing.

Twenty years of Vermont IT services taught us one thing: businesses want to know what they're getting and what it costs before they sign anything. These are our three packages.

FIXED FEE

AI Readiness Assessment

$7,500

One-time · Delivered in 2 weeks

- Audit of your current data assets
- AI opportunity map — ranked by ROI
- Written data-readiness report
- Privacy & compliance risk review
- Recommended AI stack for your industry
- 30-min findings presentation with your team

Best for: Businesses not yet sure where to start. Applies as a credit toward any build engagement.

MOST POPULAR

Custom Private AI Build

Starting at $35,000

Project-based · Scoped after assessment

- Everything in the Assessment
- Private model trained on your data only
- Custom AI agents for your key roles

- Secure deployment (on-premise or private cloud)
- Full staff training — hands-on, not a video
- 90-day post-launch support included
- IP ownership transfers to you at project close

Best for: Businesses ready to deploy AI that actually knows their operations.

RETAINER

Private AI Operations

$3,500 / mo

Monthly · Min 6-month term

- Quarterly model retraining on new data
- Usage analytics & performance reporting
- Direct access to your VAS team (not a ticket queue)
- New employee AI onboarding
- Compliance audit support (HIPAA, SOC 2)
- Priority response SLA

Best for: Post-build clients who want their AI to stay current and their team supported.

SECTION 5 — WHO THIS IS FOR

Private AI matters most when the data is sensitive, the stakes are high, and generic tools aren’t cutting it.

&– Law Firms

Document Review · Brief Drafting · Privileged Data

AI trained on your case files, briefs, and precedents — so associates draft faster and partners review less. Client privilege stays intact.

&™ Regional Manufacturers

SOPs & Manuals · Quote Automation · Knowledge Retention

AI that knows your parts catalog, supplier contracts, and production SOPs — cutting quoting time, reducing errors, and keeping institutional knowledge when employees retire.

&• Healthcare Practices

HIPAA Compliant · Clinical Notes · Zero PHI Exposure

HIPAA-ready AI for clinical documentation, patient intake, and internal operations — without any PHI touching a public system.

Ø=Ü¼ Financial Services

Audit Trails · Regulatory Docs · Access Controls

AI for client reporting, regulatory documentation, and internal research — with the audit trails and access controls your compliance team actually wants to see.

NOT IN ONE OF THESE VERTICALS?

We’ve also worked with professional services firms, logistics companies, and multi-location retail. If you have proprietary data and a real process, we can probably help. Book a discovery call and we’ll tell you directly whether private AI makes sense for your situation.

WHEN PRIVATE AI IS NOT THE RIGHT TOOL

We'll tell you if your situation doesn't warrant the investment. If public AI tools with proper data governance policies are sufficient for your use case, we'll say so. We make money on successful engagements, not on convincing you that you need something you don't.

The discovery call is free. If private AI isn't right for you, we'll tell you why and point you toward a better solution.

SECTION 6 — THE VENDOR EVALUATION CHECKLIST

5 questions to ask before you sign anything.

Use this checklist in every AI vendor conversation. If they can’t answer these questions directly, keep looking.

1. Where does the model actually run?

Look for: on-premise, private cloud, or dedicated VPC. Avoid: any answer that involves "our infrastructure" or "a secure cloud environment" without specifics. Ask for the data center location and who controls access.

2. Can you give me written documentation of what happens to my data?

Look for: data processing agreement (DPA), clear retention limits, explicit statement that data is never used for model training. Avoid: "We follow industry best practices" — that is not a legal commitment.

3. Who owns the trained model when the engagement ends?

Look for: explicit written statement that you own the model weights, fine-tuning data, and deployment configuration. Avoid: licensing arrangements, "the model remains in our cloud," or IP assignment clauses buried in terms of service.

4. How do you handle data access during the build?

Look for: scoped access, activity logging, and a defined process for revoking access at project close. Avoid: vendors who claim they "don’t access customer data" during fine-tuning — that is technically inaccurate.

5. What happens if the vendor is acquired or goes out of business?

Look for: contract clause that requires your data and models to be transferred to you in full in either scenario. Avoid: silence on this topic. Most vendors have not thought about it. We have.

THE PRIVACY CHECKLIST (COPY THIS INTO YOUR VENDOR RFIS)

'

Data never transmitted to a public LLM API in the training or inference pipeline

'

Model runs on infrastructure exclusively controlled by [CLIENT NAME]

'

Fine-tuning data, model weights, and deployment config transferred to client at project close

'

Access during build phase scoped, logged, and revoked at project close

'

Written DPA with specific retention limits (not "industry standard")

- No use of client data for model training, benchmarking, or improvement
- Right to audit data handling practices with reasonable notice
- Clear exit clause: all client assets transferred on contract termination

READY TO MOVE FORWARD?

You have three paths. Choose the one that fits.

OPTION 1

Talk to us directly

Book a free 30-minute discovery call. We'll tell you whether private AI makes sense for your business, what it would cost, and what the data looks like. No pitch, no commitment.

vermont-ai-systems.com/discovery

OPTION 2

Start with an assessment

Not ready to commit to a full build? The AI Readiness Assessment is a fixed-fee deliverable — $7,500, 2 weeks, written findings. It applies as a credit toward any build engagement.

Start with assessment !'

OPTION 3

Keep learning on your own

If you want to read more before talking, the rest of this guide covers the engagement model, pricing, and vendor evaluation in detail. Bookmark it. Share it with your IT team.

Read the full guide above

ABOUT VERMONT AI SYSTEMS

We are a Vermont-based firm that builds private, custom-trained AI for businesses that cannot expose data to public LLMs. We have worked with law firms, manufacturers, healthcare practices, and financial services companies across New England and the Northeast. We have been in business for over 20 years doing IT services — this is the natural evolution of that work.

We are not a startup. We are not venture-funded. We have no interest in becoming a large AI company. We work with a small number of clients at a time and we make sure every engagement gets the full attention of our team.

Vermont AI Systems

Burlington, Vermont

info@vermontaisystems.com

vermont-ai-systems.com

vermont-ai-systems.com/discovery