The Gittielabs Sovereign AI Readiness Audit

Moving from “Shadow AI” to Secure Business Intelligence

Are you building assets, or leaking them?

Most organizations are stuck in the "Novelty Phase" of Artificial Intelligence. Employees are using public tools (ChatGPT, Claude) to summarize emails or write code, often pasting sensitive IP into the cloud without oversight. This is "Shadow AI."

To move from ad-hoc usage to **Operational Advantage**, you must treat AI as infrastructure, not a utility.

How to use this scorecard: Select the option that best matches your organization’s current reality for each of the 5 pillars. Be honest—the goal is to identify gaps in your privacy, security, and automation architecture.

The Assessment

1. The Perimeter: How do employees access AI?

Security starts at the front door. If you don't control the interface, you don't control the data.

- [0 pts] The Wild West: Employees use personal accounts or public websites (OpenAI, Anthropic) directly. We have no visibility into who is using what.
- [3 pts] Policy Only: We have "Enterprise" licenses (e.g., ChatGPT Team), but we rely on a written policy to stop employees from sharing sensitive data.
- [5 pts] The Wrapper: We use a custom internal interface (a "Wrapper") that sits between our staff and the models. It automatically logs usage and prevents direct data leaks.

2. Data Sovereignty: Where does the data live?

Regulated industries and Government cannot afford to train public models with their secrets.

- [0 pts] Public Cloud: We paste client names, financials, or code into public models. Our data potentially trains future versions of the AI.
- [3 pts] Zero Retention: We use API keys with "Zero Data Retention" policies, but the data still leaves our secure environment (VPC) to be processed by a third party.
- [5 pts] Sovereign: We host open-source models (Llama, Mistral) on our own private cloud or on-premise servers. Our data never leaves our perimeter.

3. Knowledge Integration: How smart is the model?

An AI that doesn't know your business is just a fancy thesaurus.

- [0 pts] Generic: The AI only knows what it learned from the public internet. It hallucinates when asked about our specific projects.
- [3 pts] Manual Context: We have to manually copy-paste our SOPs or documents into the chat window every time we want an answer.
- [5 pts] RAG (Retrieval Augmented Generation): The AI is connected to a Vector Database containing our internal documents. It can cite our own data in its answers.

4. Agentic Workflow: What does the AI actually do?

Moving from "Chatting" to "Doing."

- [0 pts] Summary Only: We use AI to summarize text, write emails, or reformat code. It is a passive tool.
- [3 pts] Rules-Based: We use rigid automations (Zapier/Make). If "X" happens, the bot does "Y." If the input changes slightly, the automation breaks.
- [5 pts] Agentic: We deploy Autonomous Agents. The AI analyzes the intent of a request, plans a multi-step workflow, and executes actions (e.g., "Draft a response to this RFP and check it against our compliance PDF").

5. Governance & Oversight: The Human in the Loop

Trust, but verify.

- [0 pts] Blind: We have no logs. If an IP leak happened today, we wouldn't know until it was too late.
- [3 pts] Reactive: We can request logs from our vendor if there is an incident, but we don't monitor them actively.
- [5 pts] Redacted & Logged: We use an automated "Redaction Layer" that flags PII (SSNs, Gov Data) before it leaves the browser. We have a dashboard showing exactly how AI is impacting our KPIs.

Scoring

Add up your points: _____ / 25

0 – 10 Points: The Danger Zone

Status: High Risk / Low ROI. Your organization is exposed. You are likely leaking IP, and your "automation" is disjointed. You are using AI as a toy, not a tool.

Immediate Action: You need a Risk & Privacy Audit. Stop the bleeding before you scale.

11 – 19 Points: The Architecture Gap

Status: Transitioning. You understand the value, but your infrastructure is brittle. You are likely overpaying for APIs or relying too heavily on "Human-in-the-Loop" for safety.

Immediate Action: Deploy the Wrapper. Secure your front door and start piloting your first autonomous agent.

20 – 25 Points: Sovereign Leader

Status: Ready to Scale. You have the infrastructure to support deep automation. You are ready to move from single agents to "Agent Swarms" and custom model fine-tuning.

Immediate Action: Enterprise Tuning. Let's train a model specifically on your industry dialect.

Don’t Just Predict the Future. Architect It.

The difference between a liability and an asset is Architecture. At Gittielabs, we don't just write strategy decks. We build the Wrappers, the Private Clouds, and the Agentic Frameworks that power secure organizations.

Ready to fix your score? Book your technical consultation with our Principal Architect.

gittielabs.com Intelligence, Applied.