FRAMEWORK | IT

The Startup Leader's AI Playbook

Your guide to safe, governed, productive AI adoption for founders and operators

Table of Contents

CHAPTER 01

Executive Summary

What You'll Learn

- Why startups are the natural home for AI, and the risk hiding in how your team already uses it
- The runway, hiring, and enterprise-trust pressures AI directly addresses
- How unmanaged Shadow AI threatens your code, customer data, and next deal
- A clear path from ad hoc AI use to a safe, governed, productive program

AI is already inside your startup. Your team is already using it, whether you have approved it or not.

That is not a problem to ban. It is an advantage to harness, if you do it on purpose. Startups are the natural home for AI. You have no legacy systems to rip out, no committee to convince, and a survival incentive that bigger companies will never feel. The same tools that let a 12-person team punch like a team of 40 are the tools your competitors are reaching for right now.

Here is the tension every founder lives with. The window to build an AI advantage is open and it is narrowing. At the same time, the way your team uses AI today is quietly creating risk: customer data pasted into free chatbots, source code dropped into tools that may train on it, financial models run through platforms with no data protection guarantee. The intentions are good. The exposure is real. And the gap between using AI and using AI well is where most startups stall.

This playbook closes that gap.

The Pressures Stacking Up

The pressures on you are specific. Running out of cash drives roughly 38% to 40% of startup failures, and median runway has compressed to around 12 months, so every hour and every dollar your team recovers matters more than it would at a larger company. Hiring is the top bottleneck for most founders, and senior talent takes months to land, so doing more with the team you have is not a nice-to-have. And the moment you chase your first real enterprise deal, a security questionnaire shows up asking how you govern AI and whether you hold a SOC 2 (Service Organization Control 2) report. A single stalled enterprise deal can cost more than your entire compliance program.

The Math on the Upside

If AI saves one person 30 minutes a day on routine work, that is 10 hours a month, or 120 hours a year. Across a 12-person team, that is 1,440 hours recovered annually. At a blended cost of $65 an hour, that is roughly $93,600 in capacity reclaimed from a single use case. Now multiply that across 5 or 10 use cases. For a startup counting months of runway, that is real oxygen.

This playbook is a founder's and operator's guide to making AI work inside your company: safely, with structure, and with measurable results. You will learn what AI actually does, where to point it first, how to

keep customer data and intellectual property protected, and how to build the governance that enterprise buyers and investors now expect.

Framework IT has spent more than 16 years managing technology for small and midsized businesses across the United States. We have watched every major technology shift, from cloud migration to cybersecurity to unified communications. AI is the next one, and it is moving faster than anything before it. We built Managed Framework AI because we saw our clients facing the same set of problems: they knew AI mattered, but they did not know where to start, who to trust, or how to do it without creating new risk.

This playbook gives you the starting point. Managed Framework AI gives you the partner to execute it.

CHAPTER 02

About the Author

Adam Barney

President and Managing Partner, Framework IT

Adam Barney is President and Managing Partner of Framework IT, a premier managed IT and telecommunications firm based in Chicago. He holds a Bachelor of Science in Finance and Business Administration from the University of Illinois Urbana-Champaign, where he graduated Summa Cum Laude. With more than 15 years of executive experience in managed services and telecommunications, Adam leads with a core philosophy that technology should be user-friendly and approachable, empowering businesses to thrive in their respective industries.

Since assuming the presidency in April 2022, Adam has led a team of over 40 professionals spanning sales, information technology, operations, marketing, human resources, and fulfillment. Under his leadership, Framework IT remains committed to its inverted-pyramid approach, which ensures clients' needs and aspirations are always the company's top priority.

> ***"Clients' needs and aspirations are always the company's top priority."***

Founded in 2008, Framework IT specializes in IT support, strategy, and cybersecurity for small and midsized businesses nationwide. The company's 30+ engineers act as an extension of client businesses, proactively preventing IT issues so teams have more time to focus on what truly matters. During his career, Adam has consulted over 1,000 companies, helping them transform their technology infrastructure.

In recent years, Adam has spearheaded the adoption of artificial intelligence in Framework IT's internal operations and service delivery, positioning the company at the forefront of AI-driven IT management. He has pioneered the launch of Managed Framework AI to help clients implement AI and AI-based automation in their own businesses, enabling organizations to unlock new levels of efficiency and competitive advantage. Adam is also a founding member of The Forge AI Alliance of MSPs, an alliance of managed service providers working to accelerate the adoption of AI and automation in their own companies and those of their clients.

Under Adam's leadership, Framework IT earned a spot on the Inc. 500 Fastest Growing Private Companies in America twice and the Inc. 5000 list at least 5 times over the past decade. The company ranked as one of the Best and Brightest Places to Work in Chicago for 5 consecutive years and one of the Best and Brightest Places to Work in the Nation twice in the last 5 years. Framework IT has maintained a BBB complaint-free record since 2008.

Adam's expertise has positioned him as a sought-after voice in managed services and business technology, as a speaker and panelist at industry events. His insights have appeared in the Harvard Business Review, the Washington Post, and Fox 32 Chicago.

CHAPTER 03

AI Is Here. The Question Is How You Use It.

What Is AI, Really?

Strip away the hype and AI is a category of software that can process language, recognize patterns, generate content, and make decisions based on data. The most visible form today is the Large Language Model (LLM), the technology behind tools like ChatGPT, Claude, and Gemini. These models were trained on enormous amounts of text and can understand context, follow instructions, and produce human-quality writing, analysis, and code.

AI goes well beyond a chatbot. Modern AI platforms combine several capabilities in a single environment:

- **Conversational AI (Chat).** Ask questions, draft investor updates, summarize a customer call, rework positioning, analyze a cohort of usage data, and get research-backed answers in seconds. This is the capability most people meet first, and it is immediately useful.
- **Workflow automation.** Connect AI to the systems you already run (your customer relationship management or CRM tool, email, project tracker, data warehouse, and file storage) and build processes that run on their own. No engineering tickets required. A visual, drag-and-drop builder lets anyone design multi-step automations.
- **AI agents.** Purpose-built assistants trained for specific jobs: drafting a sales sequence, analyzing a profit and loss statement, coaching a demo call, writing release notes, responding to support tickets, and dozens more. Agents go beyond chat by following structured processes with built-in guardrails.
- **AI phone agents.** Voice-based AI that handles inbound and outbound calls, understands natural language, routes intelligently, and connects to your systems. Available around the clock, which matters when your team is small and your customers are everywhere.

Why This Matters Now for Startups

Startups win on speed and capital efficiency. AI compounds both. A relatively small product and engineering team can now ship what used to take many more people, and go-to-market teams can run plays that once required headcount you could not afford. The companies that build this muscle early will pull ahead, because the gains stack quarter over quarter.

The pressure is real on every side:

- **Runway is short.** With median runway around 12 months and cash the leading cause of startup death, capacity you recover with AI extends the most valuable thing you have: time.
- **Hiring is hard.** When a key hire takes 6 months to land, automating the work that does not require a human frees your team to do the work that does.
- **Enterprise buyers are demanding.** The first time you sell upmarket, procurement asks how you handle data and AI. Getting ahead of that question turns a blocker into a selling point.

The Real Risk Is Not AI. It Is Unmanaged AI.

The biggest threat is not that your team uses AI. It is that they are using it without your knowledge, approval, or oversight.

This is called Shadow AI, and it is happening in most companies today. Engineers paste proprietary source code into assistants to debug it. Founders drop the cap table or a board deck into a free tool to clean it up. Support reps feed customer records into a chatbot to draft a reply. The work gets done. The data leaves the building.

For a startup, the exposure is sharper than for a big company, because your differentiation often is the data and the code. According to LayerX's 2025 Enterprise AI and SaaS Data Security Report, nearly 40% of files employees upload to AI tools contain personal or payment-card data, and more than half of the text people paste into these tools includes corporate information. Most of it flows through personal accounts that sit completely outside any company control.

When someone pastes a customer list, a financial model, or proprietary code into a free AI tool, that data may be used to train the model. There is no visibility, no audit trail, and no recourse. The day a serious enterprise prospect or an investor asks how you govern AI, "we are not sure" is the wrong answer.

$4.4M	97%	40%	$120K
average cost of a data breach in 2025 (IBM)	of AI-related breach victims lacked AI access controls (IBM)	of files uploaded to AI tools hold personal or payment data (LayerX)	median enterprise deal risked without a SOC 2 report (SecureLeap)

Banning AI does not work, and it puts you further behind. The real move is to replace unmanaged, ungoverned AI with a structured, partner-led program. Same speed. Same productivity. Without the exposure that can sink a deal, a fundraise, or the company.

The Human-in-the-Loop Principle

AI is a force multiplier, not a replacement for judgment. The practical rule is to match the level of human review to the stakes of the task.

- **Low-risk tasks** can run with light oversight: drafting internal notes, summarizing a meeting, brainstorming feature names, cleaning up a Slack update.

- **Medium-risk tasks** need a human to review before anything ships: customer-facing emails, marketing copy, release notes, first-pass financial analysis.
- **High-risk tasks** require professional judgment and sign-off every time: pricing commitments, contract terms, security representations to a customer, anything that touches regulated data or a board decision.

Set this expectation early and your team gets the speed of AI without betting the company on an unreviewed output.

CHAPTER 04

What AI Can Do for Your Startup

AI is a set of tools your team can use today, for work they are already doing, to get better results in less time. The trick is knowing where to point it first. Below is a practical look across the functional areas of an early-stage company, with examples a founder or operator will recognize.

A note on scope: these are business and operational use cases. Where a task touches regulated data, customer security commitments, or legal terms, keep a human in the loop and treat AI as a drafting and analysis aid, not the final word.

Product and Engineering

Your engineers are your most expensive and hardest-to-hire resource. AI helps them move faster without cutting corners.

- **Coding assistance.** AI coding tools draft functions, write tests, explain unfamiliar code, and accelerate refactors. In the 2026 Supabase State of Startups survey, a majority of startups reported more than half their codebase was written with AI assistance.
- **Documentation and runbooks.** Turn tribal knowledge into clear technical documentation, onboarding guides for new engineers, and incident runbooks through structured interviews, so the context does not live in one person's head.
- **Research and architecture.** Use AI to compare approaches, summarize library tradeoffs, and pressure-test a design before you commit sprint time to it.

Go-to-Market: Sales and Marketing

Small teams have to run go-to-market like a much larger company. AI closes the gap.

- **Prospect research.** Compile verified intelligence on a target account and its decision-makers in minutes, including company details, leadership, recent news, and relevant context.
- **Outreach and content.** Draft personalized outbound sequences, landing pages, blog posts, and social content grounded in real positioning rather than generic filler. Teams using AI for content commonly cut content creation time by 30% to 50%.

- **Demo coaching and pipeline.** Analyze call transcripts, score them against proven sales methodologies, and surface follow-up actions. Analyze your pipeline for at-risk and stale deals you would otherwise miss.

Customer Success and Support

When the team is small, support quality is hard to keep consistent. AI makes it repeatable.

- **Draft responses.** Generate accurate, on-brand replies to tickets and emails, grounded in your help center, with a human reviewing before send.
- **Knowledge base.** Build and maintain help articles and FAQs that deflect repeat questions and scale your support without scaling headcount.
- **After-hours coverage.** An AI phone agent can handle inbound calls, answer common questions, and route the rest, so a 10-person company never sounds like one.

Fundraising and Investor Relations

Fundraising is a second full-time job for founders. AI takes the busywork off your plate.

- **Investor updates.** Produce clear monthly update memos through a structured interview that captures your KPIs (key performance indicators), metrics, and narrative.
- **Deck and narrative.** Draft and critique pitch decks and data rooms using evidence-based design principles, then stress-test the story before you walk into the room.
- **Diligence prep.** Organize and summarize the documents investors will ask for, and rehearse the hard questions with an AI that plays the skeptic.

Finance and Runway

Cash is the scoreboard. AI gives a founder without a full finance team real visibility.

- **Financial analysis.** Upload a profit and loss statement and get variance analysis, trend identification, and anomaly detection with the rigor of a seasoned analyst.
- **Forecasting and scenarios.** Build runway models and forecasts using multiple methods, with confidence scoring, so you can see the effect of a hire, a price change, or a slower sales month before you make the call.
- **ROI cases.** Quantify the business case for any spend with structured cost-benefit analysis and complexity scoring.

People and Hiring

Hiring is the top bottleneck for most founders. AI makes a lean process faster and fairer.

- **Job descriptions and scorecards.** Create clear, inclusive job descriptions and role success profiles in minutes.

- **Resume screening.** Analyze candidates against the role, compare finalists side by side, and generate structured interview plans with scoring rubrics.
- **Onboarding.** Generate onboarding checklists and 30/60/90-day plans tailored to each role, so new hires reach productivity faster.

Founder and Operations: Buying Back Time

The founder's time is the scarcest asset in the company. AI buys it back.

- **Meetings and execution.** Run tighter leadership meetings, capture decisions and action items, and turn a messy strategy session into a clean plan.
- **Decision support.** Model high-uncertainty decisions with structured scenario analysis, and use an AI devil's advocate to find the flaw in a plan before the market does.
- **Policy and process.** Draft the operational and security policies a growing company needs, including the AI use policy and the controls that support a future SOC 2.

Measuring AI ROI

Track value in 4 categories so you can prove the gains and double down on what works:

- **Time savings.** Hours recovered per week, per person, per workflow.
- **Error reduction.** Fewer rework cycles, fewer mistakes that reach a customer.
- **Revenue acceleration.** Faster sales cycles, more pipeline coverage, quicker shipping.
- **Cost avoidance.** Work absorbed without adding a hire you would otherwise have needed.

CHAPTER 05

AI Agents: Your Digital Workforce

Beyond the Chatbot

Most people's first experience with AI is a chatbot: type a question, get an answer. That is useful, but it barely scratches the surface. The real power for a startup lives in agents and automated workflows, because they let you add capacity without adding headcount.

An AI agent is a purpose-built assistant designed for a specific task or process. Unlike a general chatbot, an agent comes pre-loaded with instructions, structure, guardrails, and domain expertise. It knows what to ask, what format to follow, what to watch for, and when to stop and bring in a human.

Think of the difference between handing a new hire a blank notepad and handing them a detailed checklist with step-by-step instructions. Both can get the job done. Only one does it the same way every time.

What Makes an Agent Different from a Chat?

Structure. An agent follows a defined process. A sales coaching agent does not just review a call. It scores performance across set categories, applies recognized methodologies (MEDDIC, SPIN, RAIN), calculates conversation metrics, and produces a coaching report with trends across calls.

Guardrails. Agents include built-in quality controls. A research agent flags confidence levels and cites sources. A policy agent answers only from the connected knowledge base and refuses to speculate.

Consistency. When 5 different people use the same agent, they get consistently structured output. Quality does not depend on who wrote the prompt.

Integration. Agents can connect to your systems. They can read from your CRM, write to your project tracker, pull from your file storage, and trigger actions in other apps, without manual copy-and-paste.

Workflows: Agents That Run Themselves

If an agent is a skilled assistant, a workflow is that assistant on autopilot. Workflows connect multiple steps into automated processes that run on a schedule, respond to a trigger, or execute on demand.

Examples a startup will recognize:

- A workflow that watches your shared inbox for new inbound leads, enriches each with company research, and drafts a tailored first reply for a human to approve.
- A workflow that pulls weekly product usage data, generates a summary of activation and churn signals, and posts it to your team channel every Monday.
- A workflow that processes inbound resumes, scores them against the job description, and surfaces the top candidates for review.

You build workflows with a visual, drag-and-drop interface. No code required. Describe what you want, connect the data, define the steps, and let it run.

The Framework IT Standard Library: 100+ Ready-to-Use Agents and Workflows

When you deploy Managed Framework AI, you do not start from a blank page. Framework IT has built a library of more than 100 agents and workflows, purpose-built for the work our clients do every day, tested and refined over years of working with small and midsized businesses. Grouped for how a startup actually operates:

- **Product and engineering:** SOP (Standard Operating Procedure) Writer for runbooks and onboarding docs, Knowledge Base Architect, Deep Research Assistant, Root Cause Analysis Expert, and architect agents that help you design and audit AI agents and workflows.

- **Go-to-market:** Prospect Research, Sales Coaching and Assessment, Sales Plan Architect, Sales Funnel Guru, Blog Writing Guru, LinkedIn Viral Post Generator, Marketing Plan Expert, Webinar Campaign Builder, Marketing Image Architect, and Product Launch Plan.
- **Customer success and support:** Customer Service Coach, Email Sentiment Analyzer and Responder, FAQ Guru, and Review Response Generator.
- **Fundraising and finance:** Investor/Board Relations Aid, Presentation Content Creator, P&L Financial Analyst, FP&A Forecasting Assistant, and Use Case Builder and ROI Analysis.
- **People and hiring:** Job Description Specialist, Hiring Analyst Pro, Success Profile Builder, Professional Development Plan, and HR Company Policy Assistant.
- **Operations and governance:** Project Plan Architect, EOS L10 Meeting Assistant, Change Management Advisor, Vendor Assessment Assistant, IT Policy Creator and Editor, and Business Continuity and Disaster Recovery Plan.

These 100+ tools are available on day 1, in addition to the 500+ community-built apps already on the platform. Your team gets value immediately while you build toward custom solutions tailored to your own processes.

CHAPTER 06

The Major AI Models and Providers

One Platform, Many Models

One of the most common mistakes companies make when adopting AI is locking into a single provider. They sign up for one tool and assume they have covered AI. In reality, different models have different strengths, and the landscape changes fast.

Managed Framework AI, powered by Hatz AI, gives your company access to dozens of Large Language Models from the leading providers through a single secure platform. No separate subscriptions. No vendor lock-in. When a new model launches, it is added automatically. For a startup, that means you are never stranded on last quarter's model or paying for five subscriptions to cover your bases.

The Major Providers

OpenAI builds the GPT family, including the GPT-5 series. These are the most widely recognized models and excel at general-purpose writing, analysis, coding, and creative work. Lighter GPT tiers offer a cost-efficient option for high-volume, simpler tasks.

Anthropic builds the Claude family, including Claude Sonnet and Claude Opus models. Claude is known for strong reasoning, careful handling of complex instructions, and coding, which is why many startup

engineering teams have adopted it. It is well-suited to long-form content, detailed research, and multi-step processes.

Google builds the Gemini family, including Gemini Pro and Flash models. Google brings strong multi-modal capability (handling text, images, and code together) and a cost-efficient Flash tier for high-volume work.

Meta builds the Llama family of open models, offering strong performance at lower cost.

xAI builds the Grok models, with strong reasoning and conversational ability.

Amazon offers the Nova models, cost-efficient options for high-volume, routine tasks where speed matters most.

How to Choose the Right Model

You do not need to become a model expert. The practical guidance is simple:

- **High-stakes work** (a customer-facing deliverable, a fundraising narrative, complex reasoning): use a frontier model from OpenAI, Anthropic, or Google. They cost more per query but produce the highest quality.
- **High-volume work** (routine drafting, simple question answering, formatting): use a value model or a faster tier. They are far cheaper and quick, which keeps your credit consumption efficient while runway is tight.
- **Image generation:** the platform includes built-in image generation for product mockups, marketing assets, and pitch visuals, in both standard and premium quality.

A good rule of thumb: start with a frontier model to get quality right, then test whether a cheaper model produces comparable results. Often it does, especially for simpler tasks.

Let the Platform Pick for You: Auto LLM

If your team would rather not think about model selection at all, Managed Framework AI includes an Auto LLM selector that does the picking automatically. Auto LLM analyzes every message in real time and routes it to the most capable and cost-efficient model based on task type, complexity, files attached, and tools involved. It operates per-message, so a quick support reply and a detailed architecture review in the same conversation each get the model best suited to them. As new models join the platform, the routing improves on its own.

Auto LLM is available in chat, agents, workflows, and apps. It can be set as a personal default by any individual user, or an administrator can set it as the default for the entire company so every team member gets governed, sensible model selection out of the gate.

Auto LLM offers 3 modes so you can match the routing to the work:

- **Lite.** Fast and lightweight for everyday questions and quick back-and-forth.

- **Performance.** The balanced default. Matches the right level of model to each task: depth when you need it, efficiency when you don't.
- **Turbo.** Maximum intelligence for complex, high-stakes work where you want the platform reaching for its strongest models (a fundraising narrative, a thorny architecture decision, a board-level financial analysis).

For most startup teams, Performance mode is the right starting point. Founders and engineers can move to Turbo for high-stakes analytical work. High-volume roles like support and content can use Lite to keep credit consumption efficient on routine work.

Built-in Image Generation

Managed Framework AI includes image generation for creating custom visuals from a text description in seconds. A fast, cost-efficient tier handles product mockups and marketing images with precise text rendering, and a premium tier produces high-resolution output with legible stylized text for pitch decks, diagrams, and launch assets. No design hire or stock-photo subscription required.

CHAPTER 07

Prompt Engineering: Getting Better Results from AI

The Skill That Multiplies Everything Else

AI is only as good as the instructions you give it. A vague prompt produces a vague answer. A specific, well-structured prompt produces output that is genuinely useful, sometimes startlingly so. Prompt engineering is the practice of writing effective instructions, and it is a communication skill more than a technical one.

The Fundamentals

Be specific about what you want. Instead of "write me an email," try "write a 3-paragraph follow-up to a prospect who took our product demo last week. Professional but conversational. Reference their interest in our reporting feature and end with a clear ask to start a 2-week pilot."

Assign a role. AI performs better when you tell it who to be. "You are a seed-stage SaaS founder writing a monthly investor update" produces dramatically different output than a bare request.

Provide examples. If you want a specific format or voice, show it. Paste a past investor update you liked, a release-note style you want matched, or a positioning doc that captures your brand. AI is excellent at pattern matching.

Break complex tasks into steps. Instead of "write our go-to-market plan," walk it through: first, define the target segment. Then 3 core messages. Then the channel plan. Then a 90-day content calendar.

Iterate, do not start over. Your first prompt rarely lands perfectly. Refine with follow-ups: make the tone more direct, add a section on pricing, cut each paragraph to 2 sentences. Iteration beats starting from scratch.

Advanced Techniques

Chain of thought. Ask the AI to show its reasoning before the final answer. This produces more thoughtful output and makes errors easier to catch.

Few-shot prompting. Provide 2 or 3 examples of the input-output pattern you want, then ask it to follow the pattern. Powerful for standardizing formats across a team.

Constraint-based prompting. Tell the AI what not to do. No jargon. Under 200 words. No unverifiable claims. Constraints sharpen output.

Template prompts. Build reusable prompts for the work you repeat: investor updates, release notes, outbound sequences, support replies. Save them so your best prompt becomes a team asset, not one person's trick.

Hallucinations and AI Accuracy

AI can state something false with complete confidence. For a startup, the dangerous cases are specific: a fabricated metric in an investor update, a made-up capability promised in a sales call, or an invented citation in a security questionnaire. Treat AI output as a strong draft, not a verified fact.

Reducing the risk:

- Ground the AI in your real data by giving it the source documents to work from.
- Ask it to cite where each claim comes from.
- Verify any number, name, date, or commitment before it leaves the building.
- Keep a human in the loop on anything customer-facing or regulated.

Evaluating AI Output Quality

Before you use an output, run it through 5 questions: Is it accurate? Is it complete? Is it in the right voice? Is it appropriate for the audience? Would I be comfortable if a customer or investor saw how it was made? If the answer to any is no, refine and rerun. Aim for the 80% draft, the version that saves you most of the work and leaves the final judgment to you.

CHAPTER 08

Earning Enterprise Trust: SOC 2, Privacy, and Protecting Your IP

Why This Chapter Matters More for Startups

For an early company, security and governance are not a compliance chore. They are a growth lever and a survival issue. The day you sell upmarket, an enterprise buyer's security team decides whether you are safe to do business with. The day you raise, an investor's diligence reviews how you handle data and AI. And every day in between, your intellectual property, your source code, your models, your customer data, is the thing that makes you worth more than your competitors. Lose control of it and you lose the company's value.

This chapter covers the governance framework your startup needs and the security architecture that Managed Framework AI provides.

The Shadow AI Problem in a Startup

Shadow AI is almost certainly already happening on your team: people using ChatGPT, Gemini, Claude, or Copilot through personal accounts, without approval or oversight. The risk is not curiosity. It is what they put into those tools.

For a startup, the protected data at risk is specific and high-value: proprietary source code and models, customer records and personal data, the cap table and financials, unreleased product plans, and confidential terms from customer contracts. According to IBM's 2025 Cost of a Data Breach Report, breaches involving shadow AI were significantly more likely to result in the loss of intellectual property, and 97% of organizations that suffered an AI-related breach lacked proper AI access controls. The report also found that 63% of organizations have no AI governance policy at all.

When a buyer's chief information security officer or an investor's diligence team asks how you govern AI, you need a documented, defensible answer. Not a shrug. A policy, a platform, and a paper trail.

SOC 2: The Trust Standard That Gates Enterprise Revenue

SOC 2 (Service Organization Control 2) is a security attestation framework from the AICPA (American Institute of Certified Public Accountants) that evaluates how a company protects customer data. For business-to-business startups, it has become the baseline expectation for enterprise buyers running security due diligence.

It is built on 5 Trust Services Criteria: Security, Availability, Processing Integrity, Confidentiality, and Privacy. Security is the only mandatory one, and most startups begin there, then add others as customers require.

A few things every founder should know:

- **There are 2 report types.** Type 1 proves your controls are designed correctly at a point in time and can often be achieved in 3 to 4 months. Type 2 proves they operated effectively over a period of several months and is the long-term standard enterprise buyers prefer.
- **Timing beats panic.** The right time to start is 3 to 6 months before your first enterprise deal, not after you lose one.
- **AI usage is now in scope.** How your team uses AI, what data goes into which tools, who has access, and whether there is an audit trail, is exactly the kind of control a SOC 2 examines and a buyer's questionnaire probes. Ungoverned AI undermines the report you are working to earn.

The cost of getting this wrong is concrete. Per SecureLeap's 2026 analysis, the median enterprise deal a startup puts at risk by not having a SOC 2 report is roughly $120,000, more than the entire first-year cost of the program in many cases.

Data Privacy: Obligations That Apply Regardless of Size

Privacy law does not wait for you to reach scale. The GDPR (General Data Protection Regulation) applies to any company that serves or monitors people in the European Union, based on where the person is, not where you are. In the United States, the CCPA (California Consumer Privacy Act) and its CPRA (California Privacy Rights Act) amendments apply once you cross revenue or data-volume thresholds, and more than 20 states now have their own privacy laws. Regulators have shown they will penalize small companies, because they treat compliance as non-negotiable regardless of size.

The practical obligations that matter for AI: know what personal data you hold and where it lives, sign a Data Processing Agreement (DPA) with every vendor that touches it (including your AI platform), give people the access and deletion rights the law requires, and never let personal data flow into a tool you do not control.

Compliance-to-Platform Mapping

The table below maps the trust and privacy obligations a growing startup faces to the specific Managed Framework AI platform capabilities that address them.

Obligation	What It Requires	How Managed Framework AI Addresses It
SOC 2 Security controls	Access control, monitoring, encryption, incident response	SOC 2 Type I and II certified platform with enterprise-grade controls
AI access governance	Role-based access, usage visibility, audit trail	Role-based permissions and full, searchable usage logs
Data privacy (GDPR/CCPA)	DPA in place, data protection, subject rights	Data Processing Agreement available; data isolation and encryption
IP and source-code protection	Prevent proprietary data leaking into public models	Contractual zero-training guarantee across all model providers
Vendor due diligence	Documented evidence for buyer questionnaires	Certifications and reports that satisfy enterprise security reviews
Investor diligence readiness	Demonstrable AI governance and controls	Documented policy, platform, and audit trail in one place

The 6 Pillars of AI Governance

1. Policy. A written AI use policy defining approved tools, prohibited data types, and acceptable use. It does not need to be long. It needs to exist, be communicated, and be enforceable.

2. Access control. Role-based permissions so the right people have the right access. A new contractor should not have the same access as your chief technology officer.

3. Data protection. Contractual zero-data-training guarantees and Data Processing Agreements with every AI provider. If it is not in writing, it does not exist.

4. Auditability. Full, searchable logs of who used AI, what they entered, and what came back. This is what a SOC 2 auditor, a cyber insurer, and an enterprise buyer all want to see.

5. Training. Ongoing education so the policy is understood and followed, not quietly worked around by well-meaning teammates.

6. Customer Data and IP Stewardship. The pillar unique to a startup. Your code, models, customer data, and confidential plans are the company's value. Treat protecting them as a first-class governance responsibility, with explicit rules for what may and may not enter any AI tool, because a single leak can compromise a deal, a fundraise, or your differentiation.

The Security Architecture Behind Managed Framework AI

Managed Framework AI is built on Hatz AI, a platform that has achieved SOC 2 Type I, SOC 2 Type II, and SOC 3 certifications. Independent auditors have verified enterprise-grade controls across infrastructure, application security, access management, encryption, monitoring, and incident response. For a startup, that means you inherit a security posture you could not build yourself this early, and you can point a buyer's security team to it.

Your data is never used for AI model training. Managed Framework AI holds contractual agreements with all model providers that explicitly prohibit training on your data. Your prompt is processed and returned. It is not stored by the provider, cannot improve their models, and cannot appear in another user's results. This is a contractual guarantee, not a policy suggestion, and it is exactly what protects your source code and IP.

Encryption. Stored data is encrypted with AES-256 (the standard used by financial institutions). Data in transit uses TLS 1.2+ encryption. Keys are managed separately and rotated regularly.

Data isolation. Each company's data is logically isolated and never commingled with another customer's. Administrators have full visibility into usage, workflows, credit consumption, and audit trails.

Compliance. The platform supports GDPR, CCPA, HIPAA (with a Business Associate Agreement), and SOX (Sarbanes-Oxley Act) requirements. DPAs and BAAs are available on request, and the certifications satisfy most enterprise security questionnaires.

24/7 monitoring. Security monitoring with intrusion and anomaly detection, an incident response plan, and regular third-party penetration testing.

When a buyer or investor asks how you govern AI, this is the documented answer that turns a blocker into a point of trust.

CHAPTER 09

Your AI Adoption Roadmap: The Crawl-Walk-Run Framework

From Zero to Measurable ROI

The companies that succeed with AI follow a deliberate methodology, one that meets the team where it is, builds confidence through early wins, and expands capability over time. Gartner research shows 80% of AI licenses go unused without structured adoption support. Crawl-Walk-Run makes sure your investment produces results, not shelfware.

Why Crawl-Walk-Run?

- **Reduces risk.** Resolve governance and security questions before AI reaches everyone, not after something goes wrong.
- **Builds confidence.** Give people time to learn the basics before tackling something complex.
- **Delivers measurable wins early.** Target a high-value, low-complexity pilot to create proof that funds further investment.
- **Creates internal expertise.** Each phase develops champions and power users, so you are never solely dependent on outside help.
- **Scales sustainably.** By the Run phase, the governance, habits, and skills are already in place.

Crawl-Walk-Run is about going smart, not slow. For a startup, smart still moves fast, just with a foundation under it.

Phase 1: Crawl - Building the Foundation

Timeline: Weeks 1-4. Before anyone builds a workflow, answer the basics: who can use AI and under what rules, how access is governed, and where people go for help. This phase feels administrative, but skipping it is what creates the exposure that kills a deal later.

- **Establish an AI Acceptable Use Policy.** Define data handling, privacy, access, and the boundaries of acceptable use. For a startup chasing enterprise customers, write it with SOC 2 in mind from day 1, so the policy you create now supports the report you will need later.
- **Define roles, permissions, and user structure.** Decide who uses the platform and what they can do: administrators, builders, general users, and chat-only users. Organize people into logical groups.
- **Set the data rules that protect your IP.** Make explicit what may and may not enter an AI tool: no proprietary source code or customer data in unapproved tools, ever. This single rule prevents the most common startup exposure.

- **Sign your Data Processing Agreement and confirm zero-training.** Get the DPA and the contractual no-training guarantee in place before real data flows.
- **Onboard users and deploy general-purpose apps.** Give everyone a few immediately useful tools: writing, summarizing, research, brainstorming.
- **Identify and develop AI champions.** Pick a few enthusiastic, influential people and invest in their skills early.
- **Set an AI-forward tone from the top.** Founders use the tools visibly and talk about it. Culture change does not happen by memo, especially in a company that takes its cues from its founders.
- **Launch recurring support sessions.** Framework IT's Office Hours give the team live, low-pressure help and surface the use cases that inform the Walk phase.

By the end of Crawl: governance in place, data rules set, DPA signed, team onboarded and exploring, champions developing, and leadership visibly backing the effort.

Phase 2: Walk - Guided Exploration and First Wins

Timeline: Months 1-4. Now AI moves from curiosity to capability. The team tackles real workflows with guidance.

- **Review and deploy pre-built applications.** Deploy the agents and workflows from the 100+ Standard Library that fit your needs.
- **Customize for specific teams.** Tailor tools to product, go-to-market, support, and operations through short working sessions.
- **Pick your first pilot workflow.** Choose one that is high value, low complexity, visible, and measurable. Strong startup starting points: inbound lead enrichment and first-reply drafting, support ticket drafting, or the weekly usage and metrics summary. Document the current state before you build.
- **Build the solution with guidance.** Develop the prompt, workflow, or agent collaboratively. The people who build the pilot gain skills they carry forward.
- **Test with a small group.** Deploy to one team first, surface edge cases, and gather feedback before a wider rollout.
- **Measure ROI and document the win.** Compare against your baseline. Share it with the team and, if useful, with investors. This first documented win turns skeptics into supporters.
- **Keep a living use case repository.** Capture every use case, what was built, what it saved, and what you learned.

By the end of Walk: specialized apps deployed, first pilot automation completed and measured, ROI documented, and a repeatable process established.

Phase 3: Run - Scaling, Independence, and Continuous Improvement

Timeline: Months 4-6+. The team is now driving, with guardrails still in place.

- **Develop power users.** Invest in your strongest users as departmental AI leads who form a distributed support network.
- **Explore custom integrations.** Connect AI to your CRM, data warehouse, product analytics, and support tools to unlock far more powerful workflows.
- **Formalize business cases.** Move from informal estimates to structured analysis of time, cost, revenue, and complexity for each new initiative.
- **Enable proactive discovery.** Each team keeps a short list of automation candidates. The shift from tell me what to automate to here is what we want next is the clearest sign of a mature AI culture.
- **Build custom applications.** Your team can now build agents and workflows that solve problems unique to your company.
- **Execute broader rollouts.** Extend successful solutions across the company with clear communication and onboarding.
- **Track value and adoption.** Monitor active users, time saved, and satisfaction, and report regularly.
- **Run a continuous improvement loop.** A recurring forum reviews metrics, prioritizes new use cases, and keeps governance current as you scale and as the SOC 2 and privacy bar rises with each new enterprise customer.

By the end of Run (which never truly ends): your startup builds and deploys AI independently, tracks measurable value across functions, and keeps finding new opportunities. AI is no longer a project. It is how you work.

The People Side of AI Adoption

Even in a small company, adoption is a change-management exercise. Name the why, involve the team in choosing the first use cases, celebrate early wins publicly, and make it safe to experiment and to say what is not working. People adopt tools they helped shape.

Data Readiness: Garbage In, Garbage Out

AI is only as good as the data you point it at. Keep your customer records, documentation, and financials reasonably clean and organized, decide what is the source of truth, and the quality of every AI output improves.

CHAPTER 10

Managed Framework AI: The Complete AI Adoption Program

Not a Software Subscription. A Managed AI Adoption Program.

Managed Framework AI is a full managed AI adoption program that combines the most powerful multi-model AI platform available with a proven adoption methodology, structured training, and the ongoing guided support that turns platform access into measurable business outcomes. You get enterprise-grade tools, governance, and a named Framework IT team accountable for your results, which for a startup is like adding an AI function you could not afford to hire.

Three Pillars: Safe. Governed. Productive.

Safe. We help you adopt AI in a way that reduces risk around security, data exposure, and misuse, so leadership can approve AI with confidence and the team can use it without creating invisible liability.

Governed. We give founders and operators visibility, control, and guardrails so AI is managed on purpose. When a buyer, investor, or insurer asks how is AI governed here, you have a documented, defensible answer.

Productive. We turn AI into practical outcomes through enablement, use-case alignment, and measurable gains, so the investment compounds instead of gathering dust.

What Is Included

Access to Dozens of AI Models. ChatGPT, Claude, Gemini, Llama, and more. Never locked into one vendor, with automatic access to new models as they launch.

Unlimited Users. No per-seat pricing surprises as you grow. Every team member included.

100+ Framework IT Standard Agents and Workflows. Purpose-built tools across product, go-to-market, support, finance, and operations. Ready on day 1, plus 500+ community apps.

No-Code Workflow Builder. Build automations with drag-and-drop, 50+ native integrations, plus thousands more through Zapier. No developers pulled off the roadmap.

AI Phone Agent (ADEL). Voice AI for call handling: consistent, governed, around the clock.

AI Champion Certification. A structured training curriculum (roughly 90 minutes) that takes your team from first-time users to confident power users.

SOC 2 Type I and II Certified. CPA-audited over 6+ months. Independently certified. Your data is never used to train any public AI model, contractually guaranteed.

Office Hours: 3 Sessions Per Week

Most AI platforms hand you a login and wish you luck. Managed Framework AI includes 3 live Office Hours sessions per week, one of the most valuable parts of the program for a lean team.

Live Coaching and Q&A. Bring your questions, get answers in real time. No question is too basic. The fastest way to get unstuck.

Use Case Workshops. Structured sessions where we build a specific workflow or agent together, live, on a problem that matters to your business. Bring your own use case, like lead enrichment or support drafting, and leave with something that saves hours a week.

Show and Tell and Peer Learning. Your team shares wins and tips with each other, facilitated by ours. The best ideas spread organically.

Office Hours are low-pressure and welcoming to beginners. They accelerate individual skill and surface the use cases that shape your roadmap.

Monthly AI Strategic Business Review (AI SBR)

Every month, your Framework IT team runs a Strategic Business Review (SBR) dedicated to your AI program. We review adoption metrics, assess progress against your Crawl-Walk-Run roadmap, identify new use cases, refine the roadmap, and plan training. The SBR keeps AI moving forward as an ongoing initiative with a named partner accountable for results.

The Framework AI Resources Hub

Every Managed Framework AI client gets access to the Framework AI Resources Hub, a dedicated web portal that puts every training material, guide, tool, and enablement resource your team needs in one place. Everything is available from day 1 and updated on an ongoing basis as the platform evolves and your team's needs grow.

What you'll find in the Hub:

- **Platform Overview.** A summary of Managed Framework AI and what is included in each service plan.
- **Adoption Program Overview.** How Framework IT works with you on AI implementation and adoption, including the implementation methodology and examples of strategic AI roadmaps.
- **Getting Started Guide for Leaders.** A step-by-step resource that helps founders, CEOs, and operations leaders sponsor, launch, and sustain AI adoption across the company.

- **Getting Started Guide for Employees.** A practical guide for individual users to get up and running quickly, covering first steps, best practices, and common use cases.
- **Training Videos.** On-demand video walkthroughs of platform features and best practices, available on Framework IT's YouTube channel.
- **AI Tips and Best Practices.** Practical guidance for getting strong AI results without overspending on credits, covering prompt techniques, model selection, and everyday efficiency.
- **Training Calendar and Office Hours Schedule.** The full calendar for upcoming Office Hours sessions and any special training events Framework IT is running.
- **Industry-Specific AI Playbooks.** Tailored AI implementation strategies and use cases for specific industries, including this startup edition along with editions for legal, financial services, accounting, consulting, and more.
- **ROI Tracking Resources.** Templates and guidance for measuring and documenting AI's business impact across your organization.
- **Standard AI Tools, Workflows, and Agents.** A complete library of the standard AI tools, workflows, and agents Framework IT has built and shares with clients.
- **Industry Use Case Ideas.** Curated AI use case concepts organized by industry to help your team identify where AI can create the most value in your specific field.
- **AI Blogs and Articles.** Links to recent Framework IT blog posts and articles related to AI, covering new developments, strategies, and practical guidance.

Access the Hub at https://infohub.helpgetairight.com (password: helloframeworkit). The Hub is the central reference point your AI Champions and team members return to throughout the program.

Accelerator Plans: Go Deeper, Build Faster

For startups that want to accelerate, Framework IT offers Accelerator Plans: packages of consulting hours where our AI team works directly with yours to develop custom agents, workflows, and automations. The key word is with. We build together so your people develop the expertise to keep building on their own. The goal is not a permanent dependency. It is to get your team to the point where it can identify opportunities, architect solutions, and deploy them independently. Typical engagements include custom agent development, workflow automation that connects AI to your systems, function-specific rollouts, and advanced use cases for teams in the Run phase.

How Managed Framework AI Compares

Feature	Microsoft Copilot	ChatGPT Team	Managed Framework AI
Users Included	25 seats	25 seats	Unlimited
AI Models	1 model	1 model	Dozens of models
Adoption Support	None included	None included	Crawl-Walk-Run framework
Data Training	May use your data	May use your data	Zero training, guaranteed
Contract	Annual	Annual	Month-to-month available
Dedicated IT Partner	No	No	Named Framework IT team
Ongoing Enablement	No	No	Office Hours 3x/week, monthly SBR
SOC 2 Certified Platform	Varies	Varies	SOC 2 Type I and II
IP and Source-Code Protection	Limited	Limited	Contractual zero-training across all models
Pre-Built Agents	Limited	Limited	100+ Framework IT + 500+ community

The Transformation

Before Managed Framework AI: The team is using AI tools nobody approved. Leadership has no visibility into what is being exposed. The pressure to do something with AI grows without a plan, and the first enterprise security questionnaire is a scramble. There is risk and quiet anxiety.

After Managed Framework AI: The company has a named partner, a documented roadmap, and a trained team using governed AI every day. Output is consistent and trusted. Workflows automate what used to take hours. Leadership can show adoption metrics and prove ROI, and the answer to how do you govern AI is ready before a buyer asks. The anxiety is replaced by confidence and momentum.

The platform is Hatz AI. The methodology is Crawl-Walk-Run. The difference is Framework IT.

CHAPTER 11

AI Governance Readiness Checklist

Is Your Startup Ready?

Use this checklist to assess where you stand today. If you cannot confidently check every box, you have gaps to address before scaling AI. Score yourself honestly. Most startups start with fewer than half checked, and that is normal.

Policy and Leadership

- ☐ We have a written AI use policy that defines approved tools and prohibited data types
- ☐ Our founders and leadership have formally endorsed our AI strategy
- ☐ We have designated an internal AI Champion to lead adoption
- ☐ Our AI policy has been communicated to the whole team
- ☐ Team members have acknowledged receipt and understanding of the AI policy
- ☐ Our AI policy is reviewed and updated at least annually
- ☐ Founders actively use AI tools and visibly support the initiative

Customer Trust, SOC 2, and IP Protection

- ☐ We know which Trust Services Criteria apply to our product and have a SOC 2 target date tied to our enterprise sales timeline
- ☐ Our AI use is built to support, not undermine, a SOC 2 report (access control, logging, data rules)
- ☐ Proprietary source code and models are explicitly prohibited from unapproved AI tools
- ☐ Customer data is explicitly prohibited from unapproved AI tools
- ☐ We can produce evidence of our AI governance for an enterprise security questionnaire
- ☐ We can answer an investor's AI and data-governance diligence questions with documentation
- ☐ We have a process to review and approve new AI tools before anyone adopts them

Data Protection and Privacy

- ☐ We have contractual zero-data-training guarantees from our AI providers
- ☐ We know which AI tools our team is currently using, including free and personal-account tools
- ☐ We have a Data Processing Agreement in place with every AI vendor
- ☐ We know whether GDPR, CCPA, or other state privacy laws apply to us, and have addressed them
- ☐ We have a Business Associate Agreement in place if we handle protected health information
- ☐ Team members know which types of data they can and cannot enter into AI tools

Access Control and Permissions

- ☐ AI access is role-based, not everyone has the same permissions
- ☐ We maintain searchable logs of AI usage (who, what, when)
- ☐ We can produce an AI governance report if asked by an auditor, insurer, or buyer
- ☐ User permissions are reviewed and updated when roles change
- ☐ Administrative access is restricted to authorized personnel only
- ☐ We have an offboarding process that includes revoking AI platform access

Training and Adoption

- ☐ Our team has received formal AI training, not just a policy memo
- ☐ We have shared prompt templates and best practices documented
- ☐ We are tracking AI adoption metrics (usage, time saved, ROI)
- ☐ We have a structured plan to expand AI usage over the next 6 months
- ☐ New hires receive AI onboarding as part of standard orientation
- ☐ We have identified and developed AI Champions within the team
- ☐ We maintain a use case repository documenting AI wins and lessons learned
- ☐ Our team has access to recurring support (Office Hours, coaching, help desk)

Workflow and Automation Maturity

- ☐ We have identified our top 3 to 5 high-value AI use cases
- ☐ We have documented baseline metrics (time, cost, error rates) for at least 1 target workflow
- ☐ We have completed at least 1 AI pilot and measured results
- ☐ We have a pipeline of future automation opportunities
- ☐ We have mapped which pre-built agents and workflows align with our highest-value use cases
- ☐ Our AI automations include human review checkpoints for critical outputs

Governance and Continuous Improvement

- ☐ We have a recurring forum to review AI governance
- ☐ We review AI adoption metrics and ROI data at least monthly
- ☐ We have a process for prioritizing new AI use cases based on business impact
- ☐ We regularly share AI wins and best practices across the team
- ☐ We have a feedback loop between users and whoever manages the AI program
- ☐ Our AI roadmap is a living document, updated as we progress and as the enterprise trust bar rises

How to Read Your Results

0-15 boxes checked: You are in the early stages. Most startups start here. The Crawl phase of the Crawl-Walk-Run framework is built exactly for this.

16-30 boxes checked: Meaningful progress, real gaps remaining. You are likely ready for the Walk phase: targeted use cases while you shore up governance and training.

31-40 boxes checked: A mature posture. You are in or approaching the Run phase, ready to scale, build custom solutions, and drive continuous improvement.

41+ boxes checked: A high level of maturity. Focus on continuous improvement, broader rollouts, and deeper automations.

If you have unchecked boxes, you are not alone. Most startups do. Managed Framework AI was built to help you check every one, with a named partner who handles the platform, the governance framework, the training, and the ongoing enablement so you can adopt AI with confidence.

CHAPTER 12

Next Steps

You have the playbook. The question now is what you do with it.

If your startup is ready to move from unmanaged AI usage to a structured, governed, productive AI program, here is how to start:

Request a consultation. We will walk through where you stand today, identify the highest-value opportunities, and show you exactly how Managed Framework AI works, including how it helps you stand up the AI governance enterprise buyers and investors look for. No pressure, no pitch deck. Just a conversation about what makes sense for your company.

Take the readiness checklist to your team. Use it to start an internal conversation about AI governance, risk, and opportunity. The checklist alone is often enough to surface gaps nobody was talking about.

Stop the bleeding on Shadow AI. Every day your team uses unmanaged AI tools is another day your source code, customer data, and IP are at risk, and another day closer to an enterprise questionnaire you are not ready for. Getting onto a governed platform is the single highest-impact step you can take, and it does not require a long procurement cycle.

Framework IT has spent more than 16 years helping small and midsized businesses get technology right. AI is the next chapter, and we are here to help you write it.

Book a Consultation

Framework IT
www.frameworkit.com
(312) 564-5446