



---

# The Consulting Firm Leader's AI Playbook

Your guide to safe, governed, productive  
AI adoption for management and strategy consulting firms

## Table of Contents

|           |  |
|-----------|--|
| <b>01</b> | Executive Summary                                      |
| <b>02</b> | About the Author                                       |
| <b>03</b> | AI Is Here. The Question Is How You Use It.            |
| <b>04</b> | What AI Can Do for Your Firm                           |
| <b>05</b> | AI Agents: Your Digital Workforce                      |
| <b>06</b> | The Major AI Models and Providers                      |
| <b>07</b> | Prompt Engineering: Getting Better Results from AI     |
| <b>08</b> | AI Security, Privacy, and Governance                   |
| <b>09</b> | Your AI Adoption Roadmap: The Crawl-Walk-Run Framework |
| <b>10</b> | Managed Framework AI: The Complete AI Adoption Program |
| <b>11</b> | AI Governance Readiness Checklist: Is Your Firm Ready? |
| <b>12</b> | Next Steps   |

---

**CHAPTER 01**

# Executive Summary

---

## What You'll Learn

- Why AI is reshaping consulting firms from the bottom up
- The compliance composite: confidentiality, E&O, and client pass-through
- The cost of shadow AI inside small to mid-sized firms
- How a governed program turns pressure into leverage

AI isn't coming to consulting. It's already inside your firm, and it's quietly reshaping what consulting work looks like.

According to the Management Consultancies Association's January 2026 survey, 77% of consulting firms have integrated AI into their systems or enabled employees to use AI models, far outpacing the 16% business average. Harvard Business Review reports that AI is restructuring the consulting firm itself, automating tasks traditionally handled by junior consultants and forcing partners to rethink staffing models, pricing, and the very definition of leverage. McKinsey's 2025 layoffs were a warning shot: when AI absorbs the work, the pyramid changes shape.

The story inside small to mid-sized consulting firms is messier than the headline number suggests. Adoption is happening, but it's happening from the bottom up. Associates are pasting client interview notes into free chatbots. Engagement managers are running competitor research through consumer AI tools. Business development teams are drafting proposals using platforms that may train on the very content that distinguishes your firm. The intentions are good. The exposure is enormous.

For consulting firm leaders, the stakes go beyond efficiency. Your firm runs on three assets: client confidentiality, the quality of your advice, and your professional reputation. AI used carelessly damages all three at once.

The pressures stacking up on consulting firms aren't going away:

- More than half of the U.S. workforce reports burnout, and SHRM data puts the consulting industry's rate above 45%. Long hours, unrealistic utilization expectations, and excessive administrative load are the consistent culprits.
- Clients expect faster turnarounds, lower fees, and the same depth of insight. AI-enabled competitors are using the productivity gain to undercut your pricing and accelerate your timelines.
- Errors and Omissions (E&O) insurance carriers have begun excluding AI-related liabilities from policies. Verisk's standardized AI exclusion forms took effect January 1, 2026, and Berkley, AIG, and Great American are following. Coverage now depends on how your firm governs AI, not just whether you use it.
- Talent markets have loosened, but the work hasn't gotten easier. Firms are doing more with smaller teams, and the pressure on every billable hour is rising.

Consider the math. If AI saves a single consultant 45 minutes per day on research, drafting, and synthesis, that's 15 hours per month, or 180 hours per year. At a blended cost of \$125 per hour (conservative for management consulting professionals), that's \$22,500 per person annually. For a firm with 25 consultants, that's \$562,500 recovered from a single use case. Now multiply that across business development, deliverable production, knowledge management, and practice operations.

But the real prize isn't the savings. It's the leverage. AI lets your firm pursue more proposals, take on more engagements, and deliver higher-quality work without proportionally expanding headcount. That's the difference between a consulting firm that compounds and one that runs in place.

This playbook exists to close the gap between "our team is already using AI" and "our firm governs AI with the same rigor we apply to every other professional standard." Over the following chapters, you'll learn what AI can do for your firm today, how to protect client confidentiality and your E&O posture, and how to build a phased adoption plan that satisfies your professional obligations while compounding operational advantages.

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

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 mid-sized 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.

# 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 of AI today is the Large Language Model (LLM), the technology behind tools like ChatGPT, Claude, and Gemini. These models were trained on massive amounts of text and learned to predict the most likely next word in a sequence. That simple mechanism, repeated billions of times across enormous datasets, produces something that looks like understanding: the ability to follow instructions, reason through problems, and produce human-quality writing, analysis, and code.

But AI goes well beyond a chatbot. Modern AI platforms combine multiple capabilities into a single environment:

- **Conversational AI (Chat).** Ask questions, draft client memos, summarize interview transcripts, analyze market data, brainstorm engagement structures, and get research-backed answers in seconds. This is the capability most people encounter first, and it's immediately useful for consulting professionals.
- **Workflow automation.** Connect AI to your existing business systems (your CRM, email, project management tools, file storage, and more) and build automated processes that run on their own. No coding required. A visual, drag-and-drop builder lets anyone design multi-step automations, from pulling client status updates into a weekly partner brief to routing inbound RFPs to the right practice lead.
- **AI agents.** Purpose-built AI assistants trained for specific tasks: synthesizing research, drafting RFP responses, generating client status reports, coaching consultants on sales conversations, 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 calls intelligently, and integrates with your business systems. Available around the clock.

## Why This Matters Now

The window for competitive advantage is open, but it's narrowing. Firms that build AI capability now will compound those gains over the next 2 to 3 years. Those that wait will face a steeper learning curve, higher costs, and competitors who have already baked AI into their proposal cycles and delivery models.

The numbers from the front lines are hard to ignore:

- 77% of consulting firms have integrated AI or enabled employee usage, compared with a 16% business average (Management Consultancies Association, January 2026).

- AI-powered proposal platforms have produced 55% to 80% time savings on construction and infrastructure proposals, with comparable wins documented across management consulting.
- BlackRock's research team increased throughput 5x through AI integration. The same pattern is playing out at small and mid-sized consulting firms applying AI to research and synthesis.
- 71% of organizations regularly use generative AI in at least one business function (Deloitte, 2024), and consulting firms cite faster data processing and deeper client insights as the top advantages.

Consider the math. If AI saves a consultant 45 minutes per day on research, drafting, and synthesis, that's 15 hours per month, or 180 hours per year. At a blended cost of \$125 per hour, that's \$22,500 per person annually. For a firm with 25 consultants, \$562,500 recovered from a single use case. The numbers compound quickly across multiple use cases.


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

The biggest threat to your firm isn't that AI doesn't work. It's that your team is already using it without your knowledge, approval, or oversight.

This is called Shadow AI, and it's happening inside consulting firms right now. Associates are pasting client interview transcripts, internal financial models, and proprietary analytical frameworks into free chatbots. Engagement managers are running competitor research and benchmarking data through consumer AI tools with no data protection guarantees. Business development teams are drafting RFP responses on platforms that may use your firm's intellectual property to train their models.

The intentions are good. The risk is enormous.

When someone pastes a client deliverable, a board memo, an interview transcript, or a proprietary methodology into a free AI tool, that data may be used to train the model. There's no visibility, no audit trail, and no legal recourse. For a firm that holds confidential client information, takes on professional responsibility for its advice, and depends on its intellectual property for competitive differentiation, this isn't a theoretical risk. It's an existential one.



The shadow AI premium isn't theoretical. IBM's 2025 Cost of a Data Breach Report found that breaches involving shadow AI take 247 days to detect on average and expose data across multiple environments 62% of the time. 83% of organizations lack the technical controls to detect or prevent employees from uploading confidential data to AI platforms.

In consulting, the consequences hit faster and harder. A leaked client deliverable is a contractual breach. A hallucinated regulatory citation in an advisory memo is professional negligence. A proprietary methodology fed to a public model is intellectual property dilution. And a published Stanford study documents the case of a consulting firm that produced a report using Azure OpenAI containing

non-existent references and fabricated court quotes, leading to corrections, a partial refund, and reputational damage.

Banning AI doesn't work, and it puts your firm further behind. The real move is replacing unmanaged, ungoverned AI usage with a structured, partner-led program. Same productivity. Zero exposure.

## The Human-in-the-Loop Principle

AI is a tool for augmenting human judgment, not replacing it. In consulting, where every recommendation carries the weight of professional advice and contractual responsibility, maintaining human oversight is nonnegotiable.

Think of AI outputs on a risk spectrum:

- **Low risk: internal research and drafting.** Using AI to summarize industry reports, draft internal meeting agendas, structure brainstorms, or compile preliminary competitive intelligence. These tasks benefit from AI speed without creating external exposure.
- **Medium risk: analysis that informs client work.** Synthesizing interview notes, benchmarking data, drafting deliverable outlines, generating proposal sections. AI accelerates the work, but a qualified consultant reviews every output before it reaches the client.
- **High risk: client deliverables, advisory memos, regulatory-adjacent guidance.** Any output that reaches a client, drives a major decision, or touches a regulated domain requires careful review by an experienced consultant. AI produces the draft. A professional owns the final product.

The principle is simple: AI handles the volume. Your team handles the judgment. The combination is what creates the advantage, and it's what keeps your professional liability posture defensible.

## CHAPTER 04

# What AI Can Do for Your Firm

---

AI is a set of tools your teams can use today, for work they're already doing, to get better results in less time. The key is knowing where to point it.

Below is a practical look at how AI applies across core consulting functions. Every use case described here focuses on operational, administrative, and analytical support work. AI doesn't replace the consultant's judgment or the professional responsibility that comes with client advice. It frees up more of your team's time for the work that actually requires their expertise.

***Important: When your firm advises clients in regulated industries (healthcare, financial services, legal, etc.), your AI usage inherits their compliance obligations. The use cases below are administrative and analytical support functions. Anything that touches a client's regulated data needs to be governed by the framework outlined in Chapter 8.***

## Business Development and Proposals

Business development consumes time that could otherwise be billable, and the work is repetitive enough that AI handles much of it well.

- **RFP response drafting.** AI can generate first drafts of RFP responses by pulling from your firm's prior responses, capability statements, case studies, and methodology documents. What used to take 3 days of senior consultant time becomes a 1-day review and refinement cycle.
- **Proposal customization.** Tailor existing proposal templates to a specific prospect's industry, business model, and stated priorities. AI matches your firm's voice while incorporating client-specific framing.
- **Prospect research.** Compile verified intelligence on target companies and their leadership in minutes, including company details, key executives, recent news, strategic priorities, and likely pain points. This sharpens every outbound conversation.
- **Pursuit qualification.** Run quick assessments of inbound opportunities against your firm's ideal client profile, win history, and capacity, helping partners decide which opportunities deserve full pursuit.
- **Pitch material development.** Draft and refine pitch decks, capability overviews, and credentials documents for new business meetings, grounded in your firm's actual experience and positioning.

## Client Delivery and Research

This is where AI delivers the most visible productivity gain. Research, synthesis, and deliverable drafting are exactly what large language models do well.

- **Deep research with cited sources.** AI agents can produce structured research reports on industries, competitors, regulatory environments, and market dynamics with cited sources in hours rather than days. This is how junior consultant time gets multiplied.
- **Interview synthesis.** Upload transcripts from client interviews and get structured summaries, theme extraction, quote pulls, and cross-interview pattern identification. What used to be 6 hours of analyst work becomes 30 minutes of review.
- **Document analysis.** AI can process strategic plans, financial reports, market studies, and operational documents, extracting key findings, flagging inconsistencies, and generating structured summaries that highlight what matters for your engagement.
- **Financial analysis.** Upload a client's P&L and get expert-level variance analysis, trend identification, anomaly detection, and scenario modeling. AI processes financial data with the rigor of an experienced analyst.
- **Deliverable drafting.** Generate first drafts of executive summaries, recommendation sections, implementation plans, and supporting analyses based on your engagement materials, methodology, and prior work.

- **Scenario planning.** Build structured scenario analyses for strategic decisions, complete with probability assessments, signposts, and recommended responses for each scenario.

## Knowledge Management and Intellectual Property

Consulting firms run on accumulated knowledge, but most of it sits trapped in folders, Slack channels, and the heads of senior partners. AI surfaces it.

- **Case study retrieval.** Find relevant prior engagements, methodologies, and deliverables instantly. A new associate gains access to 10 years of firm expertise on day 1.
- **Methodology codification.** Turn senior consultants' tacit knowledge into documented frameworks through structured interviews. What used to take a 6-month internal project takes weeks.
- **Internal Q&A.** Maintain a master FAQ as a single source of truth for client situations, engagement approaches, and firm policies, with confidence-scored answers grounded in your actual knowledge base.
- **SOP creation.** Turn process knowledge into professional, audit-ready Standard Operating Procedures through structured interviews. Engagement playbooks, onboarding processes, and quality review workflows become formal assets instead of tribal knowledge.

## Client Communications and Reporting

Client communication is high-frequency, high-stakes, and disproportionately consumes senior consultant time. AI accelerates the routine parts.

- **Status report drafting.** Generate first drafts of weekly engagement updates, monthly status reports, and quarterly stakeholder briefings using engagement data and prior reports.
- **Board and executive memos.** Draft structured update memos for client board members or executive sponsors through guided interviews, with style matching to your firm's voice.
- **Email management.** Analyze inbound client emails for sentiment and urgency, then draft response options that match the appropriate tone and capture action items.
- **Meeting prep.** Generate structured meeting prep documents, agendas, and pre-read materials grounded in engagement context.
- **Presentation development.** Create polished, brand-compliant slide content from engagement materials, with the design work handled by AI image generation when needed.

## Practice Operations

Running a consulting practice is itself a complex operation, and AI handles much of the administrative load.

- **Project planning.** Build engagement plans with work breakdown structures, timelines, RACI matrices, and risk registers in hours rather than days.
- **Resource and utilization planning.** Pull data from time tracking and project systems, identify utilization gaps, surface overcommitted team members, and flag projects running behind schedule.
- **Engagement quality review.** Run preliminary quality reviews on draft deliverables for clarity, completeness, alignment to scope, and adherence to firm standards before partner review.
- **Risk and issue management.** Track engagement risks, conduct root cause analysis when issues surface, and develop corrective action plans through structured AI-guided workflows.
- **Vendor and subcontractor management.** Evaluate technology vendors, subcontractor candidates, and partnership opportunities with structured assessment frameworks.

## Finance, Billing, and Utilization

Consulting firms live and die by utilization, realization, and cash flow. AI provides analytical horsepower that small firms previously couldn't justify hiring for.

- **Financial analysis and forecasting.** Get expert-level P&L analysis, scenario modeling, and forecasting for your firm's own financials, with auto-method assignment and confidence scoring.
- **Pricing and proposal economics.** Model engagement profitability under different staffing scenarios, pricing structures, and timeline assumptions. Sharpen go/no-go decisions.
- **Utilization reporting.** Aggregate utilization data, identify trends across practices, and flag patterns that need partner attention.
- **Billing review.** Run preliminary reviews of pre-bills, time entries, and expense reports for completeness and consistency.
- **AI ROI analysis.** Quantify the productivity gains from your own AI program with structured ROI calculations that hold up to partner scrutiny.

## Human Resources and Talent

Talent is the constraint in consulting. AI is not a substitute for great people, but it changes how you find, develop, and retain them.

- **Job description creation.** Generate inclusive, effective job descriptions for consultant, analyst, and operations roles in minutes.
- **Resume screening.** Analyze candidate resumes against role requirements, surface multi-candidate comparisons, and flag potential interview questions or concerns.

- **Performance management.** Coach managers through Professional Development Plans, performance reviews, and improvement plans with structured, evidence-based methodologies.
- **Onboarding.** Generate comprehensive onboarding plans for new consultants, including curriculum design, mentorship structures, and first-engagement assignments.
- **Sales coaching for senior consultants.** Analyze recorded client conversations and provide developmental coaching, scorecards, and conversation metrics for consultants who are still building their business development muscle.

## Measuring AI ROI

ROI from AI falls into 4 categories. When you're building the case internally or evaluating results, measure against all 4:

### Time savings

Hours reclaimed per person per week on tasks AI now handles or accelerates. This is the most immediately visible category and translates directly to billable capacity.

### Error reduction

Fewer mistakes in research synthesis, deliverable drafting, financial analysis, and client communications. Errors in consulting work are expensive to fix and damaging to client trust.

### Revenue acceleration

Faster proposal turnaround, quicker engagement starts, and more responsive client communications that compress sales cycles and accelerate cash collection.

### Cost avoidance

Confidentiality incidents prevented, shadow AI exposure eliminated, and headcount growth avoided by doing more with existing teams. For consulting firms, this category also includes professional liability exposure reduction.

# AI Agents: Your Digital Workforce

---

## Beyond the Chatbot

Most people's first experience with AI is a chatbot: you type a question, you get an answer. That's useful, but it barely scratches the surface. The real power of enterprise AI lives in agents and automated workflows.

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 questions to ask, what format to follow, what mistakes to watch for, and when to stop and ask for human input.

Think of the difference between handing a new analyst a blank notepad and handing them a detailed engagement workplan with step-by-step instructions and quality checks. Both can produce work. Only one does it consistently, every time, without missing steps.

## What Makes an Agent Different from a Chat?

**Structure.** An agent follows a defined process. A financial analysis agent, for example, doesn't just "review these numbers." It performs variance analysis, identifies anomalies, runs scenario models, generates forecasts, and produces a downloadable report with trend tracking.

**Guardrails.** Agents include built-in quality controls. A research agent will flag confidence levels, cite sources, and explicitly note when it's operating outside its reliable knowledge. A policy agent will only answer from the connected knowledge base and will refuse to speculate.

**Consistency.** When 10 different consultants use the same agent, they get consistently structured outputs. The quality doesn't depend on who wrote the prompt or how senior they are. That consistency is exactly what makes a deliverable defensible and a firm scalable.

**Integration.** Agents can connect to your business systems. They can read from your CRM, write to your project management tools, pull data from your file storage, and trigger actions in other applications, all without manual copy-and-paste.

## Workflows: Agents That Run Themselves

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

Examples relevant to consulting firms:

- A workflow that monitors industry news, regulatory updates, and competitor moves in your clients' sectors, scores them for relevance, and surfaces the top items to engagement teams weekly.
- A workflow that pulls weekly engagement status data from project management tools, generates structured client updates, and routes them to engagement managers for review.
- A workflow that processes inbound RFP notifications, scores them against your firm's pursuit criteria, drafts initial qualification notes, and routes promising opportunities to the right partner.

You build workflows using a visual, drag-and-drop interface. No coding required. You describe what you want, connect the data sources, 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 don't start from scratch. Framework IT has built a library of over 100 agents and workflows, purpose-built for the types of work our clients do every day. These are tools we've designed, tested, and refined based on years of working with small and mid-sized businesses.

The library spans functions directly relevant to consulting firm operations:

- **Business development and proposals.** RFP builder, RFP vendor Q&A generator, prospect research, B2B marketing letters, sales funnel analytics, webinar campaign builder, presentation content creator.
- **Client delivery and research.** Deep research assistant, scenario planning strategist, P&L financial analyst, FP&A forecasting, use case builder and ROI analysis, brainstorming partner, devil's advocate for stress-testing recommendations.
- **Knowledge management.** Knowledge base architect, FAQ guru, SOP writer, sales playbook curator, trend and topic researcher.
- **Client communications and reporting.** Investor and board relations aid, leader email structure enforcer, email sentiment analyzer, presentation slide generator, biography writer.
- **Practice operations.** Project plan architect, vendor assessment assistant, root cause analysis expert, business continuity and disaster recovery plan, IT policy creator.
- **Finance and utilization.** P&L financial analyst, FP&A forecasting assistant, AI ROI break-even calculator, Excel expert.

- **HR and talent.** Job description specialist, hiring analyst pro, success profile builder, professional development plan coach, performance improvement plan, change management advisor, sales coaching and assessment, career coach.
- **Productivity.** Prompt coach, SMART goal setter, EOS L10 meeting assistant, Microsoft Word and Excel support.

These 100+ tools are available on day 1 of your deployment. They are in addition to the 500+ community-built apps already available on the platform. Your team can start getting value immediately while you build toward custom solutions tailored to your firm's specific methodologies.

## CHAPTER 06

# The Major AI Models and Providers

---

## One Platform, 69 Models

One of the most common mistakes firms make when adopting AI is locking themselves into a single provider. They sign up for ChatGPT and assume they've "covered AI." In reality, different AI models have different strengths, and the landscape changes fast.

Managed Framework AI, powered by Hatz AI, gives your firm access to 69 Large Language Models (LLMs) from the leading AI providers, all through a single secure platform. No separate subscriptions. No vendor lock-in. When a new model launches, it's added to the platform automatically.

## The Major Providers

**OpenAI** builds the GPT family of models, including GPT-5, GPT-5.1, GPT-5.2, and GPT-5.4. These are the most widely recognized AI models and excel at general-purpose tasks including writing, analysis, coding, and creative work. GPT-5 Nano offers a cost-efficient option for high-volume, simpler tasks.

**Anthropic** builds the Claude family, including Claude 4.5 Sonnet, Claude 4 Opus, Claude 4.5 Opus, Claude Opus 4.6, and Claude Sonnet 4.6. Claude models are known for strong reasoning, nuanced analysis, and careful handling of complex instructions. They are particularly well-suited for long-form content, detailed research, and tasks that require following multi-step processes, exactly the kind of work consulting deliverables require.

**Google** builds the Gemini family, including Gemini 2.5 Pro, Gemini 3 Pro, Gemini 3.1 Pro Preview, and Gemini Flash 3. Google's models bring strong multi-modal capabilities (handling text, images, and code together) and tight integration with Google's search and data infrastructure. Gemini 3.1 Flash is an excellent cost-efficient option.

**Meta** builds the Llama family of open-source models, including Llama 4 Maverick and Llama 4 Scout. These models offer strong performance at lower cost and represent Meta's commitment to open AI development.

**xAI** builds Grok 3 and Grok 4, high-capability models with strong reasoning and conversational abilities.

**Amazon** offers Nova Micro and Nova Lite, cost-efficient models well-suited for high-volume, routine tasks where speed matters more than peak capability.

**Specialized models** include advanced reasoning models designed for mathematical problems, structured analysis, and multi-step logical reasoning, exactly the kinds of tasks that appear in financial modeling, scenario planning, and complex client deliverables.

## How to Choose the Right Model

You don't need to become an AI model expert. The practical guidance is straightforward:

- **High-stakes work** (client deliverables, advisory memos, board-facing presentations): Use frontier models from OpenAI, Anthropic, or Google. They cost more per query but produce the highest quality output.
- **High-volume work** (internal status updates, routine email drafting, preliminary research notes): Use value models from Amazon, Meta, or the faster tiers from Google and OpenAI. They are significantly cheaper and fast.
- **Long-form drafting and reasoning**: Anthropic's Claude family is particularly strong for sustained analytical writing, multi-section deliverables, and nuanced strategic recommendations.
- **Image generation**: The platform includes Nano Banana (powered by Gemini) for creating custom images from text descriptions, useful for client presentation visuals, capability deck graphics, and marketing materials.

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

## The Power of Model Flexibility

The AI model landscape changes fast. A model that leads today may be outperformed next quarter. Managed Framework AI insulates your firm from this churn. You're never locked into one vendor's technology. When new models launch, they're available on your platform automatically. No new contracts, no new subscriptions, no rebuild required.

This matters for consulting firms that need to match the right tool to the right task. Your research analysts might prefer Claude for detailed competitive analysis, use GPT-5 for drafting executive memos, and run Gemini for multi-modal slide development, all in the same platform, sometimes in the same conversation. That flexibility is a material operational advantage.

## Built-in Image Generation

Managed Framework AI includes Nano Banana, the platform's image generation capability. Describe what you want to see, and it creates a custom image in seconds. Two tiers are available:

- **Nano Banana 2** (Gemini 3.1 Flash Image): Fast, cost-efficient image generation with real-time web search knowledge, precision text rendering, and translation for client presentations and capability mockups.
- **Nano Banana Pro** (Gemini 3 Pro Image): Premium quality output with high-resolution options (up to 4K) and legible stylized text generation for client-facing infographics, executive briefing visuals, and marketing assets.

No design skills or stock photo subscriptions 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 for AI, and it's a communication skill more than a technical one. Think of it the same way you'd think about briefing a new analyst. The better the brief, the better the work product.

## The Fundamentals

**Be specific about what you want.** Instead of "Summarize this market," try "Summarize the current state of the U.S. industrial automation market, including market size, top 5 players by revenue, regulatory pressures, and emerging technology trends. Format the output as a 2-page briefing memo suitable for a private equity client evaluating a platform investment. Use a professional but direct tone."

The more context you provide, the better the output. Include details about the audience, the purpose, the format, the tone, and any constraints.

**Assign a role.** AI performs better when you tell it who to be. "You are a senior management consultant with 12 years of experience in operations transformation for industrial companies" produces dramatically different output than a bare request to "analyze this operating model."

**Provide examples.** If you want output in a specific format or style, show the AI what good looks like. Paste in a previous deliverable you were proud of, a slide structure you want replicated, or a client memo

that matches your firm's voice. AI is excellent at pattern matching, and your firm's accumulated work product is a goldmine of patterns.

**Break complex tasks into steps.** Instead of asking AI to "develop a market entry strategy," walk it through the process: "First, summarize the target market. Then identify the top 3 entry models. Then assess each against the client's capabilities. Then recommend a preferred path with a 90-day roadmap." Step-by-step instructions produce step-by-step quality.

**Iterate, don't start over.** Your first prompt rarely produces a perfect result. That's normal. Refine the output by giving follow-up instructions: "Make the tone more direct." "Add a section on competitive dynamics." "Shorten the executive summary to 3 paragraphs." Iteration is faster than starting from scratch.

## Advanced Techniques

**Chain of thought prompting.** Ask AI to show its reasoning before giving a final answer. "Walk me through your analysis step by step before providing your recommendation." This produces more thoughtful, accurate outputs and makes it easier to spot errors in the reasoning, which matters when AI output will become the basis for client advice.

**Few-shot prompting.** Provide 2 or 3 examples of the input-output pattern you want, then ask AI to follow the pattern with new input. This is especially powerful for standardizing formats across your firm, like project status reports, board memos, or post-engagement summaries.

**Constraint-based prompting.** Tell AI what not to do. "Do not make assumptions about the client's financial position. Do not use industry jargon the board wouldn't understand. Do not exceed 500 words." Constraints sharpen the output and prevent common problems.

**Template prompts.** Build reusable prompt templates for tasks your team performs regularly. A standardized prompt for drafting weekly status updates, RFP responses, or executive briefings ensures consistent quality across the team and shortens the learning curve for new hires.

## Hallucinations and AI Accuracy

AI models sometimes generate information that sounds confident but is factually incorrect. Stanford research found that general-purpose AI tools produce inaccurate outputs at rates between 58% and 88% on professional tasks. In a consulting context, the consequences of acting on hallucinated market data, fabricated regulatory citations, or invented case studies can be severe: damaged client relationships, contractual breaches, and professional liability exposure.

The Stanford research isn't theoretical. A documented case involved a consulting firm that produced a report using Azure OpenAI containing non-existent references and fabricated court quotes, leading to corrections, a partial refund, and reputational damage. Insurance carriers are now writing AI exclusions into E&O policies precisely because of this risk.

Reducing the risk:

- **Never treat AI output as verified fact.** AI produces drafts, not gospel. Every claim, every number, every citation needs human verification before it reaches a client deliverable, a regulator, or a board.
- **Ask AI to cite its sources.** When using models with web search capability, request cited sources so you can verify independently.
- **Use AI for what it's good at.** Summarization, drafting, pattern recognition, and structured analysis are strengths. Generating new factual claims out of thin air is a weakness.
- **Build verification into your workflow.** Treat AI-generated analysis the same way you'd treat a first-year analyst's work: review it, challenge it, and confirm the numbers and citations.

## Understanding Context Windows

Every AI model has a "context window," the amount of text it can process in a single conversation. Think of it like working memory. A model with a 200,000-token context window can process roughly 150,000 words at once, enough to analyze an entire engagement workroom of key documents in a single prompt.

Practical tips for consulting professionals:

- Larger context windows allow you to upload entire client document sets, interview transcripts, and prior deliverables in a single conversation.
- Quality degrades as you approach the context limit. Put the most important content first.
- Break very large analysis tasks into focused chunks rather than dumping everything into one prompt.
- Save key AI-generated analysis to reference documents rather than relying on conversation history.

## Evaluating AI Output Quality

When AI produces a draft, ask 5 questions before using it:

- **Is it factually accurate?** Verify every specific claim, statistic, and citation.
- **Is it complete?** Did AI miss any important angle or consideration?
- **Is it appropriately specific?** Generic AI output gets you 60% there. Specific output for your client and their situation gets you 90%.
- **Is the tone right?** Does it match your firm's voice and the audience's expectations?
- **Would I be comfortable putting my name on it?** If the answer is no, keep iterating.

Treat AI as a tool that produces an 80% draft. Your judgment, expertise, and revision turn it into a 100% deliverable. That's not a bug. That's the design.

# AI Security, Privacy, and Governance

---

## Protecting Your Firm in the Age of AI

Security is the foundation everything else sits on. Every other benefit of AI, the productivity gains, the proposal acceleration, the competitive advantages, means nothing if your firm's intellectual property is exposed, your clients lose confidence in your confidentiality, or your firm faces a professional liability claim because someone pasted client information into an unsecured tool.

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

## The Shadow AI Problem in Consulting

Shadow AI is already happening inside consulting firms: team members using AI tools like ChatGPT, Gemini, Claude, and Copilot without your knowledge, approval, or oversight.

The risk isn't about employee curiosity. It's about what they're putting into those tools. When a consultant pastes interview transcripts, client financial data, proprietary methodologies, or draft deliverables into a free AI tool, that data may be used to train the model. There's no visibility, no audit trail, and no legal recourse.

For consulting firms, the data at risk is uniquely sensitive:

- **Client confidential information.** Strategic plans, financial projections, operational data, customer lists, and other information your firm is contractually obligated to protect under NDAs and Master Service Agreements (MSAs).
- **Engagement deliverables.** Reports, recommendations, models, and analyses that represent your firm's intellectual contribution and your client's strategic direction.
- **Firm intellectual property.** Methodologies, frameworks, proprietary research, and pricing models that differentiate your firm competitively.
- **Client-industry regulated data.** When your clients operate in regulated industries (healthcare, financial services, legal, defense), the data they share with you may carry compliance obligations like HIPAA, SEC, FINRA, or ITAR. Your AI usage inherits those obligations.

When your client's general counsel, your professional liability insurer, or a prospect's procurement team asks how your firm governs AI, you need a documented, auditable answer. Not a shrug. A policy, a platform, and a paper trail.

## Client Confidentiality, E&O Exposure, and Industry Pass-Through: The Compliance Landscape

Consulting firms don't operate under a single regulatory framework like medical practices under HIPAA or financial services firms under SEC rules. Instead, your compliance obligations come from three overlapping sources: your contracts, your professional liability posture, and your clients' regulatory environments.

**Client confidentiality (NDAs and MSAs).** Every engagement begins with confidentiality obligations. Master Service Agreements typically include strict prohibitions on disclosing client information, mishandling client data, or using client materials outside the engagement scope. Pasting confidential client information into a free AI tool may breach those obligations. M&A NDAs increasingly include explicit provisions prohibiting upload of confidential information to public or open-source AI platforms, and sellers are restricting the use of AI tools that retain data or use inputs for model training.

**Professional liability and E&O exposure.** Insurance carriers have responded to the rise of AI-related claims by writing AI exclusions into E&O policies. Verisk's standardized AI exclusion forms became effective January 1, 2026, providing insurers with ready-made language to exclude losses arising from generative AI use, including specific tools like ChatGPT, Midjourney, and DALL-E. Major carriers including Berkley, AIG, and Great American are filing for regulatory approval to exclude AI-related liabilities. Professional liability coverage will increasingly depend on how firms govern their AI usage, with insurers seeking detailed information on acceptable use policies, staff training, and human-in-the-loop protocols. Governance is now a coverage condition, not a nice-to-have.

**Client-industry pass-through compliance.** When you advise clients in regulated industries, your handling of their data inherits their compliance framework. Healthcare clients trigger HIPAA Business Associate obligations. Financial services clients trigger SEC, FINRA, and Reg S-P obligations. Legal clients trigger attorney-client privilege protections. Defense clients trigger ITAR and CMMC requirements. If your AI tools don't satisfy those frameworks, your firm cannot legally serve those clients on engagements that touch regulated data.

**State privacy laws.** California's CCPA/CPRA, Colorado's CPA (including its Colorado AI Act provisions), Virginia, Connecticut, and a growing list of state privacy laws apply to consulting firms that handle personal information about consumers or employees. AI tools that lack appropriate data handling controls may create compliance violations.

**International obligations.** Consulting firms with European clients or operations face GDPR obligations, including data protection impact assessments for AI processing of personal data. The EU AI Act, which entered application in August 2026 for general-purpose AI obligations, imposes additional duties on the deployment of AI systems.

**Contractual representations.** Many large client contracts now include specific AI-related representations: that you will not use AI tools without disclosure, that AI usage will comply with the client's policies, that AI output will be reviewed by qualified professionals, and that confidentiality controls extend to any AI systems your firm uses. Violating these representations is a contractual breach.

The compliance-to-platform mapping below shows how Managed Framework AI's security architecture addresses each of these requirements:

| Compliance Requirement                 | What It Demands   | How Managed Framework AI Addresses It                               |
|--|---|---|
| Client Confidentiality (NDAs/MSAs)     | Protect client data; restrict use; honor retention limits | Zero data training guarantees; data isolation; contractual DPAs     |
| Professional Liability / E&O           | Documented AI governance; human-in-the-loop protocols     | SOC 2 Type II; acceptable use policy; audit logs; role-based access |
| Industry Pass-Through (HIPAA)          | BAA available for healthcare clients                      | Business Associate Agreement available on request                   |
| Industry Pass-Through (SEC/FINRA)      | Audit trail; supervision; books and records               | Searchable audit logs; admin controls; usage history                |
| State Privacy Laws (CCPA, CPA, etc.)   | Data protection; deletion rights; transparency            | AES-256 encryption; data isolation; deletion controls               |
| EU AI Act / GDPR                       | DPIA support; high-risk system governance                 | DPAs available; documented governance framework                     |
| Vendor Due Diligence (client-mandated) | SOC 2; security documentation; subprocessor disclosure    | SOC 2 Type I and Type II; SOC 3; documentation packet               |

## The 6 Pillars of AI Governance for Consulting Firms

**1. Policy.** A written AI use policy defining approved tools, prohibited data types, and acceptable use cases. For consulting firms, this policy must explicitly address client confidential information, engagement deliverables, firm IP, and the special handling required for clients in regulated industries. Without a policy, you can't enforce your position or defend it to clients, insurers, or prospects.

**2. Access control.** Role-based permissions ensuring the right people have access to the right AI capabilities. A new analyst shouldn't have the same AI access as a managing partner. A subcontractor shouldn't have access to firm-wide tools. A well-designed permission structure prevents confusion, limits risk, and ensures appropriate guardrails from day 1.

**3. Data protection.** Contractual zero-data-training guarantees and Data Processing Agreements (DPAs) with every AI provider. If it isn't in writing, it doesn't exist. Your platform must guarantee that your firm's data, your clients' data, and your engagement materials are never used to train any public AI model. Full stop.

**4. Auditability.** Full, searchable logs of who used AI, what prompts were entered, and what was returned. This is required to satisfy contractual representations to clients, demonstrate governance to E&O insurers, and provide your clients' compliance teams with the transparency they're increasingly demanding.

**5. Training.** Ongoing education so your policies are understood and followed, not quietly worked around. A policy memo doesn't count. Real training means structured onboarding, recurring support sessions, and a culture where everyone from analysts to partners knows the rules and the reasons behind them.

**6. Client Confidentiality and Engagement Compliance.** This pillar is unique to consulting firms. Establish clear classification rules for client confidential data, engagement deliverables, and information that triggers client-industry compliance obligations. Define which data categories can enter the AI platform and under what conditions. Build handling procedures that align with your firm's NDAs, MSAs, and the regulatory frameworks your clients operate under. When a healthcare client asks whether your AI usage is HIPAA-compliant, you need to be able to say yes and prove it.

## 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 security controls across infrastructure, application security, access management, encryption, monitoring, and incident response.

**Your data is never used for AI model training.** Managed Framework AI has contractual agreements with all AI model providers (OpenAI, Anthropic, Google, Meta, xAI, Amazon, and others) that explicitly prohibit training on your data. When you submit a prompt, your data is processed and returned. It isn't stored by the provider. It can't improve their models or appear in responses to other users. This is a contractual guarantee, not a policy suggestion.

**Encryption.** All stored data is encrypted using AES-256 encryption (the same standard used by financial institutions and government agencies). All data in transit uses TLS 1.2+ encryption. Encryption keys are managed separately and rotated regularly.

**Data isolation.** Each organization's data is logically isolated. Your firm's data is never commingled with other clients' data. Administrators have full visibility into usage, workflows, credit consumption, and audit trails.

**Compliance.** The platform supports compliance with GDPR, CCPA, HIPAA (with Business Associate Agreement), and SOX. DPAs and BAAs are available upon request. The certifications satisfy most enterprise security questionnaires and provide evidence of due diligence for clients reviewing your vendor management program.

**24/7 monitoring.** Security monitoring with intrusion detection and anomaly detection, backed by an incident response plan with immediate containment, notification, root cause analysis, and remediation. Regular third-party penetration testing validates the security posture.

***When your client's general counsel asks "How do you govern AI?" you need a documented, auditable answer. A policy, a platform, and a paper trail.***

CHAPTER 09

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

## From Zero to Measurable ROI

The firms that succeed with AI follow a deliberate methodology, one that meets teams where they are, builds confidence through early wins, and progressively expands capability.

Gartner research shows 80% of AI licenses go unused without structured adoption support. The Crawl-Walk-Run framework makes sure your investment produces results, not shelfware.

## Why Crawl-Walk-Run?

Firms that rush AI deployments encounter the same problems: low adoption, confused users, security concerns, and an inability to prove value. The Crawl-Walk-Run model addresses each of these by design:

- **Reduces risk.** Resolve governance and security questions before AI reaches a broad audience, not after a client raises a concern.
- **Builds confidence.** Give people time to learn the basics before asking them to do something complex. They adopt tools more readily and with greater enthusiasm.
- **Delivers measurable wins early.** Target a high-value, low-complexity pilot workflow in the Walk phase to create tangible proof of value.
- **Creates internal expertise.** Each phase develops champions, power users, and eventually an internal community of practice.
- **Scales sustainably.** By the time you reach the Run phase, the governance, cultural habits, and technical skills are already in place.

| Phase | Timeline    | Focus      | Key Milestones   |
|-------|-------------|------------|--|
| Crawl | Weeks 1-4   | Foundation | AI policy in place, roles defined, users onboarded, client data classification completed, champions identified |
| Walk  | Months 1-4  | First wins | Pilot workflow live, ROI measured, specialized tools deployed, client confidentiality controls tested          |
| Run   | Months 4-6+ | Scale      | Custom integrations, firm-wide deployment, continuous improvement, custom methodologies digitized              |

## Phase 1: Crawl - Building the Foundation

### Timeline: Weeks 1-4

The Crawl phase lays the groundwork. Before anyone builds a workflow or automates a process, the firm needs to answer fundamental questions about governance, access, and data classification.

**1. Establish an AI Acceptable Use Policy.** Create a formal AI Acceptable Use Policy covering data handling, privacy, access controls, compliance requirements, and the boundaries of acceptable AI use. For consulting firms, this policy must explicitly address client confidential information, engagement deliverables, firm intellectual property, and the special handling required when clients operate in regulated industries.

**2. Classify your data for AI use.** Define which categories of firm and client data can enter the AI platform and under what conditions. Establish clear handling rules for each engagement based on the client's industry, the data sensitivity, and any contractual restrictions in your MSA or NDA.

**3. Define roles, permissions, and your user structure.** Partners, engagement managers, consultants, analysts, business development, operations, and administrative staff likely need different levels of access. Organize users into logical groups.

**4. Execute Data Processing Agreements.** Ensure DPAs are in place with the AI platform provider. Review the platform's SOC 2 report. Document this in your vendor management files for client due diligence requests. If you serve healthcare clients, ensure a Business Associate Agreement is in place.

**5. Review your E&O posture.** Notify your professional liability insurer of your AI program. Confirm that your governance framework satisfies their underwriting requirements. Update your engagement letter language to address AI usage if needed.

**6. Onboard users and deploy general-purpose AI applications.** Bring users onto the platform through a controlled process. Deploy a curated set of general-purpose applications: writing assistants, research tools, brainstorming aids.

**7. Identify and develop AI champions.** Select a small group of enthusiastic team members and invest in their skills early. Champions should complete the AI Champion Certification.

**8. Establish an AI-forward culture from leadership.** Partners must visibly use AI tools and actively encourage teams to experiment. Culture change doesn't happen by memo.

**9. Launch recurring support sessions.** Framework IT's Office Hours provide live, recurring sessions where users get hands-on help with real work.

**By the end of Crawl:** Governance framework in place, data classification defined, E&O posture validated, users onboarded and exploring, champions developing expertise, compliance documentation started.

## Phase 2: Walk - Guided Exploration and First Wins

**Timeline: Months 1-4**

The Walk phase is where AI moves from curiosity to capability. The firm tackles real workflows with guidance.

**1. Deploy pre-built applications to targeted teams.** Review the 100+ Standard Apps library and deploy the ones relevant to your firm's operations: deep research, P&L financial analysis, RFP builder, presentation content creation, knowledge base architect.

**2. Identify your first pilot workflow.** Working with your AI champions and Framework IT, identify a pilot that's high value, low complexity, and measurable. Strong candidates for consulting firms include RFP response drafting, deep research synthesis for active engagements, weekly status report automation, or interview transcript synthesis.

**3. Develop the AI solution with guidance.** Build the pilot as a collaborative effort. The people involved develop skills they carry into future projects.

**4. Test with a controlled pilot group.** Deploy to a small group first, typically one practice area or one engagement team. Surface edge cases and gather feedback before broader rollout.

**5. Measure ROI and document your first win.** Quantify results: hours saved per consultant per week, proposal turnaround time, deliverable cycle time, utilization impact. Share the results with the partnership.

**6. Test client confidentiality controls in practice.** Use the Walk phase to validate your data classification rules and handling procedures with real (non-critical) scenarios before scaling. Confirm that your audit trail satisfies what your largest clients would ask for in a security questionnaire.

**By the end of Walk:** Pilot workflow completed and measured, specialized tools deployed, ROI documented, confidentiality controls tested in practice.

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

**Timeline: Months 4-6+**

The Run phase marks the transition from guided experimentation to organizational self-sufficiency.

**1. Develop a power users program.** Invest in advanced training for team members who've demonstrated strong AI skills. They become the internal experts who help your firm scale capability.

**2. Explore custom integrations.** Connect AI to your practice management, CRM, time tracking, and project management tools. This is where workflow automation starts to deliver compounding returns.

**3. Digitize your proprietary methodologies.** Convert your firm's signature frameworks, playbooks, and intellectual property into structured AI agents. This is how a 25-person firm starts delivering the consistency of a 250-person firm.

**4. Automate end-to-end workflows.** Build workflows that handle proposal generation from RFP intake through final review, or weekly status reporting from project data through client delivery.

**5. Formalize business case analysis.** Move from informal estimates to formal ROI analysis for new AI initiatives. Use the same rigor on your own firm's investments that you'd apply to a client's business case.

**6. Establish a continuous improvement loop.** Create a recurring forum where partners, practice leaders, and operations review adoption metrics, prioritize new use cases, and update governance.

**By the end of Run (which never truly ends):** Your firm is building and deploying AI solutions independently, tracking measurable value, and continuously identifying new opportunities. AI is no longer a project. It's part of how your firm operates.

## The People Side of AI Adoption

Technology adoption is a people problem, not a technology problem. Consulting professionals range from AI enthusiasts to skeptics, and your adoption plan needs to account for both.

4 principles that work:

- **Start with pain, not technology.** Don't lead with "here's an AI tool." Lead with "here's something that saves you 4 hours a week on research synthesis." The tool is the how, not the why.
- **Make early adopters visible.** When an associate uses AI to cut a deliverable draft from 8 hours to 2, make sure the partnership hears about it. Wins compound when they spread.
- **Lower the bar.** Frame the first step as "ask it a question" not "build an automation." People who start with a simple question become the ones building complex workflows 3 months later.
- **Address the fear.** Some consultants worry AI will replace them, especially given recent industry restructuring. The honest answer: AI replaces tasks, not professionals. The consultants who learn to use AI will outperform those who don't, and your firm needs your best people using better tools.

## Data Readiness: Garbage In, Garbage Out

AI is only as useful as the data you feed it. For consulting firms, data readiness means:

- **Organized firm knowledge.** If every case study, every methodology, and every prior deliverable is scattered across consultants' personal drives, AI can't help you reuse them. Centralize what you can.
- **Consistent engagement records.** Status reports, deliverables, and engagement summaries should follow consistent templates so AI can learn the patterns.
- **Clean CRM data.** Your business development effort is only as good as your CRM hygiene. Before automating prospect tracking or proposal generation, clean up the underlying data.

## 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.

## Three Pillars: Safe. Governed. Productive.

**Safe.** We help consulting firms adopt AI in a way that reduces risk around client confidentiality, professional liability, and the regulatory frameworks your clients operate under, so that partners can approve AI with confidence and team members can use it without creating invisible exposure for the firm.

**Governed.** We give firm leadership greater visibility, control, standards, and guardrails so AI use is managed intentionally rather than emerging chaotically. When a client's general counsel, your E&O insurer, or a prospect's procurement team asks "How is AI governed at your firm?" you have a documented, defensible answer.

**Productive.** We turn AI into practical business outcomes through enablement, use-case alignment, workflow improvement, and measurable gains in efficiency and execution, so that the investment in AI compounds over time instead of quietly gathering dust.

## What Is Included

**Access to 69 AI Models.** ChatGPT, Claude, Gemini, Llama, DeepSeek, and more. Never locked into one vendor. Automatic access to new models as they are released.

**Unlimited Users.** No per-seat pricing surprises. Every team member included.

**100+ Framework IT Standard Agents and Workflows.** Purpose-built tools covering business development, client delivery, research, financial analysis, knowledge management, client communications, operations, HR, and productivity. Ready on day 1, in addition to 500+ community-built apps on the platform. Easily cloned and customized for your firm.

**No-Code Workflow Builder.** Build powerful automations with drag-and-drop. 58+ native integrations plus thousands more through Zapier. Zero developers needed.

**AI Phone Agent (ADEL).** Voice AI for call handling: consistent, governed, available around the clock. Useful for inbound business development and after-hours support.

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

**AI Navigator Training and Quick Tips Training.** Recurring training assets that keep your team current as the platform and models evolve.

**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.

Office Hours are working sessions, facilitated by Framework IT's AI team, where your people get hands-on help with real work. The format rotates across 3 types:

**Live Coaching and Q&A.** Bring your questions. Get answers. Whether you're struggling with a prompt for synthesizing client interview notes, trying to figure out which model to use for an executive memo, or wondering how to approach a specific workflow.

**Use Case Workshops.** Structured sessions where we build specific workflows together. We've run workshops on RFP response automation, deliverable drafting, knowledge management buildouts, and methodology digitization.

**Show and Tell and Peer Learning.** Users share wins, discoveries, and creative applications with each other. Someone in business development shows how they automated prospect research. Someone on a delivery team shares a prompt that generates 80%-complete deliverable outlines. The collective intelligence accelerates everyone's learning.

## Monthly AI Strategic Business Review (SBR)

Every month, your Framework IT team conducts a Strategic Business Review (SBR) dedicated to your AI program. This is a structured strategic session designed to keep your AI adoption on track.

During each SBR, we review adoption metrics and usage patterns, assess progress against your Crawl-Walk-Run roadmap, identify new use cases and automation opportunities, update and refine your AI roadmap, and plan training and enablement activities. The SBR keeps AI adoption moving forward with a named partner accountable for results.

## Accelerator Plans: Go Deeper, Build Faster

For firms that want to accelerate their AI maturity, Framework IT offers Accelerator Plans: packages of consulting hours where our AI team works directly with yours to develop custom agents, workflows, and automations.

Typical engagements include custom agent development for firm-specific methodologies, end-to-end proposal workflow automation, practice-specific rollouts, and advanced use case development for firms in the Run phase. The goal is to get your firm to the point where your own team can identify opportunities, build solutions, and deploy them independently.

## 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           | 69 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   |
| Pre-Built Agents                 | Limited           | Limited           | 100+ Framework IT + 500+ community  |
| Client Confidentiality Controls  | Limited           | Limited           | Data isolation, zero training, DPAs |
| E&O Governance Documentation     | No                | No                | SOC 2, audit logs, policy templates |
| Industry Pass-Through Compliance | No HIPAA BAA      | No HIPAA BAA      | HIPAA BAA available, audit-ready    |

## The Transformation

**Before Managed Framework AI:** The firm is operating in the dark. Consultants are using AI tools nobody approved. Partners have no visibility into what client information is being exposed. Your E&O insurer has started asking about AI governance, and you don't have a documented answer. Clients are starting to ask about AI policies in their procurement processes, and there's no framework to show them.

**After Managed Framework AI:** The firm has a named partner, a documented roadmap, and a trained team actively using governed AI every day. Proposals get out the door faster. Research and synthesis take a fraction of the time. Client deliverables go through AI-assisted drafting and human-led refinement.

Compliance can demonstrate AI governance to clients, insurers, and prospects on demand. The anxiety is gone, 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 Firm Ready?*

Use this checklist to assess where your firm stands today. If you can't confidently check every box, you have gaps that need to be addressed before scaling AI adoption. Score yourself honestly. Most firms start with fewer than half of these checked, and that's normal.

### Policy and Leadership

- We have a written AI use policy that defines approved tools and prohibited data types
- Our partners have formally endorsed our AI strategy
- We have designated an internal AI Champion to lead adoption
- Our AI policy has been communicated to all firm employees and subcontractors
- Employees have acknowledged receipt and understanding of the AI policy
- Our AI policy is reviewed and updated at least annually
- Firm leadership actively uses AI tools and visibly supports the initiative

### Client Confidentiality and Engagement Compliance

- Our AI policy explicitly addresses client confidential information and engagement materials
- We have reviewed our standard MSA and NDA language for AI-related provisions
- Our engagement letters address AI usage and human-in-the-loop review
- We can demonstrate AI governance compliance to clients on request
- We have identified which clients operate in regulated industries (healthcare, financial services, legal, defense)
- We have appropriate BAAs, DPAs, or other agreements in place for regulated client work
- We have a process for client-mandated AI restrictions or disclosure requirements

### Data Protection and Privacy

- We have contractual zero-data-training guarantees from our AI providers
- We know which AI tools our employees are currently using (including free tools)
- Sensitive data categories (client confidential info, engagement deliverables, firm IP) are defined as off-limits for public AI
- We have Data Processing Agreements in place with every AI vendor
- We have a defined process for classifying data sensitivity levels for AI use

### 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 a client, insurer, or auditor
- We have a process for reviewing and approving new AI tools before deployment
- User permissions are reviewed and updated when roles change
- Administrative access is restricted to authorized personnel only
- We have a defined 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 employees receive AI onboarding as part of their standard orientation
- We have identified and developed AI Champions within the firm
- We maintain a use case repository documenting AI wins and lessons learned
- Our team has access to recurring support (Office Hours, coaching, help desk)

### Professional Liability and E&O

- We can answer "How does your firm govern AI?" with a documented response
- Our E&O insurer is aware of our AI usage and governance framework
- We have reviewed our E&O policy for AI exclusions and conditions
- Our AI governance includes documented human-in-the-loop protocols
- We have a process for verifying AI-generated content before client delivery
- We have reviewed industry-specific AI regulations applicable to our clients

### Workflow and Automation Maturity

- We have identified our top 3 to 5 high-value AI use cases
- We have documented baseline metrics (time, cost, quality) 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 firm's highest-value use cases
- Our AI workflows include human review checkpoints for client-facing outputs

### Governance and Continuous Improvement

- We have a recurring forum (steering committee, review meeting) for 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 firm
- We have a feedback loop between users and the team managing the AI program
- Our AI roadmap is a living document that is updated based on progress and new opportunities

### How to Read Your Results

**0-15 boxes checked:** You are in the early stages. Most firms start here, and there's nothing wrong with that. The important thing is to start building the foundation. The Crawl phase of the Crawl-Walk-Run framework is designed exactly for this.

**16-30 boxes checked:** You have meaningful progress but significant gaps remain. You are likely ready to move into the Walk phase, tackling targeted use cases while shoring up governance and training.

**31-40 boxes checked:** Your firm has a mature AI governance posture. You are in or approaching the Run phase, ready to scale, build custom solutions, and drive continuous improvement.

**41+ boxes checked:** You are operating at a high level of AI maturity. Focus on continuous improvement, expanding to new teams, and deepening the complexity of your automations.

If you have unchecked boxes, you aren't alone. Most firms do. Managed Framework AI was built to help you check every one of them. As a managed AI adoption program, our team handles the platform, the governance framework, the training, and the ongoing enablement so you can adopt AI with confidence, with a named partner who stays accountable for results.

**CHAPTER 12**

# Next Steps

---

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

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

**Request a consultation.** We'll walk through where your firm stands today, identify the highest-value opportunities, and show you exactly how Managed Framework AI works. No pressure, no pitch deck. Just a conversation about what makes sense for your firm.

**Take the readiness checklist to your partners.** Use it to start an internal conversation about AI governance, client confidentiality exposure, and operational opportunity. The checklist alone is often enough to surface gaps that nobody was talking about.

**Stop the bleeding on Shadow AI.** Every day your team uses unmanaged AI tools is another day your client confidential information, engagement deliverables, and firm intellectual property is at risk. Getting onto a governed platform is the single highest-impact step you can take, and it doesn't require a long procurement cycle.

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

## Book a Consultation

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