



The Investment Banking Leader's AI Playbook

Your guide to safe, governed, productive
AI adoption for boutique and middle-market firms

Table of Contents

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

CHAPTER 01

Executive Summary

What You'll Learn

- Why AI adoption has become a competitive and supervisory imperative for investment banks
- How Shadow AI creates MNPI, confidentiality, and FINRA Rule 3110 exposure
- The economic math of AI productivity gains for a 25-person firm
- How this playbook closes the gap between unmanaged AI usage and a governed program

AI isn't coming to investment banking. It's already here, and it's reshaping how deals get sourced, pitched, diligenced, and executed.

More than 80% of Tier 1 investment banks have integrated AI across their front, middle, and back offices, and the largest firms are reporting hundreds of millions in productivity gains. Goldman Sachs deployed its GS AI Assistant to roughly 46,500 employees and cut pitch creation time by 50%. AI agents now process a Confidential Information Memorandum (CIM) in 30 to 60 minutes, a job that used to consume an analyst for 4 to 8 hours. The competitive gap between firms that have AI in their workflow and those that don't is widening every quarter.

For boutique and middle-market firms, the stakes go beyond efficiency. The Financial Industry Regulatory Authority (FINRA) issued Regulatory Notice 25-07 in April 2025, formally extending Rule 3110 supervisory duties to generative AI workflows. The 2026 FINRA Annual Regulatory Oversight Report identifies AI as a continuing and emerging priority, and FinCEN imposed an \$80 million Bank Secrecy Act penalty on a broker-dealer in March 2026, the largest in the agency's history. The regulatory floor is moving fast, and your firm's AI posture is now part of every exam, every audit, and every client diligence questionnaire.

Meanwhile, your analysts and associates are already using AI. They're pasting deal memos, financial models, and pitch outlines into free chatbots. Coverage bankers are drafting client communications with consumer tools that may train on the contents. Buyer lists and target screens are being run through platforms that have no contractual protection for your firm's confidentiality obligations. The intentions are good. The exposure is enormous, and it sits directly on top of material nonpublic information (MNPI), client engagement letters, and Non-Disclosure Agreement (NDA) obligations.

The Pressures Stacking Up

The pressures stacking up on small and midsized investment banks aren't going away:

- Pitch volume keeps climbing while team sizes don't. Middle-market deal volume hit a record 12,856 transactions in 2025, up 17.1% year over year. Fee pressure is rising in parallel: only 34% of M&A advisors raised fees in 2024, down from 47% the year before.
- Junior banker burnout is a structural problem. Analyst attrition runs about 85% within 2 years at major firms. Replacing an analyst costs roughly 150% of their salary, and JPMorgan estimated \$500 million

in attrition-driven losses in 2024 alone. Boutiques feel the same pressure on a smaller absolute base, with deeper relative damage.

- Hours haven't gotten shorter. The typical analyst still works 70 to 100 hours per week, and 72% report damage to family and personal relationships. The work that AI can absorb (pitch formatting, data scraping, deck building, comp pulls) is exactly the work driving people out.
- The compliance perimeter keeps expanding. FINRA Rule 3110, Rule 4511, and Rule 2210 all apply to AI-generated work, and Securities and Exchange Commission (SEC) recordkeeping rules under 17a-4 may apply to AI outputs that get transmitted externally.

The Math

Consider the math. If AI saves a single banker 30 minutes per day on routine work, that's 10 hours per month, or 120 hours per year. At a blended hourly cost of \$150 (conservative for licensed banker professionals), that's \$18,000 per person annually. For a 25-person firm, that's \$450,000 recovered from a single use case. Now multiply that across pitch automation, CIM analysis, buyer list research, due diligence, and back office. The numbers get serious fast.

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 a control room log." Over the following chapters, you'll learn what AI can do for your firm today, how to protect MNPI and client confidentiality, how to satisfy FINRA and SEC supervisory expectations, and how to build a phased adoption plan that compounds operational advantages without creating new regulatory exposure.

Framework IT has spent more than 16 years managing technology for small and mid-sized 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 sections of pitch books, summarize a 10-K, analyze a target's financial statements, brainstorm a deal thesis, and get research-backed answers in seconds. This is the capability most people encounter first, and it's immediately useful for banker work.
- **Workflow automation.** Connect AI to your existing business systems (your Customer Relationship Management platform, email, file storage, deal pipeline tools, 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 generating a weekly pipeline summary to triaging incoming CIMs by sector and stage.
- **AI agents.** Purpose-built AI assistants trained for specific tasks: drafting executive summaries, generating buyer lists, building comp tables, screening targets against an investment thesis, summarizing earnings calls, 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. Useful for centralized reception, executive assistant overflow, and after-hours coverage.

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. Firms that wait will face a steeper learning curve, higher costs, and competitors who have already baked AI into their pitch process and deal execution.

The numbers from the front lines are hard to ignore:

- More than 80% of Tier 1 investment banks have integrated AI across front, middle, and back offices, with disclosed financial impact in the low hundreds of millions of dollars per firm.

- Goldman Sachs deployed its GS AI Assistant to roughly 46,500 employees and reports a 50% reduction in pitch creation time.
- AI agents now process a CIM in 30 to 60 minutes, compared to 4 to 8 hours of manual review.
- AI-assisted target screening shrinks per-company review time from roughly 45 minutes to 8 minutes.
- Major banks report cutting deal preparation time by more than 30% using generative AI for due diligence, pitch books, and client presentations.

Consider the math. If AI saves a single banker 30 minutes per day on routine work, that's 10 hours per month, or 120 hours per year. Multiply that across a team of 25, and you're looking at 3,000 hours recovered annually. At a blended cost of \$150 per hour, that's \$450,000 in productivity gains from a single use case. Now multiply that across 5 or 10 use cases. The numbers get serious fast.

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 investment banks right now. Analysts are pasting target financials, deal memos, and pitch outlines into free chatbots. Associates are running comp tables through consumer tools with no data protection guarantees. Vice presidents are drafting CIM sections using platforms that may train on confidential client information. Coverage bankers are routing buyer outreach drafts through tools that retain prompts indefinitely.

The intentions are good. The risk is enormous.

When someone pastes a draft deck, a buyer list, a fairness opinion analysis, or an MNPI-laden deal memo 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 an investment bank that handles MNPI every day, operates under FINRA supervisory obligations, and signs engagement letters that promise client confidentiality, this is not a theoretical risk. It is a direct breach of duty.

80%+

of Tier 1 IBs have
integrated AI across the
firm

68%

surge in shadow gen AI
usage in enterprises

\$5.56M

average financial services
data breach cost

Rule 3110

now applies to AI under
FINRA Notice 25-07

Shadow AI incidents now account for roughly 20% of all data breaches and carry a \$670,000 cost premium versus standard breaches. The IBM 2025 Cost of a Data Breach Report identifies financial services as the second-most expensive industry to suffer a breach. Gartner predicts that by 2030, more than 40% of enterprises will experience security or compliance incidents tied to unauthorized AI use.

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 investment banking, where every output may reach a client, a regulator, or an investment committee, 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, build a sector primer, or compile preliminary competitive intelligence. These tasks benefit from AI speed without creating external exposure.
- **Medium risk: analysis that informs decisions.** Comp table construction, target screening, draft management presentations, draft CIM sections. AI accelerates the work, but a banker reviews every output, validates the data, and signs off before it moves forward.
- **High risk: client deliverables, communications with the public, regulatory filings.** Pitch books going to a board, fairness opinion analyses, FINRA Rule 2210-governed marketing materials, regulatory submissions. AI can produce a draft. A licensed professional owns the final product and the supervisory review.

The principle is simple: AI handles the volume. Your team handles the judgment. The combination is what creates the advantage.

CHAPTER 04

What AI Can Do for Your Firm

AI is a set of tools your team 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 investment banking functions. Every use case described here focuses on operational, administrative, and analytical work. AI doesn't replace banker judgment, client relationships, or the supervisory review your firm owes its regulators. It frees up more of your team's time for the work that actually requires it.

Important: The use cases described in this chapter are administrative and analytical. They do not constitute legal, tax, accounting, or investment advice. All client deliverables and regulatory filings require qualified professional review.

Deal Origination and Sourcing

Origination teams spend a disproportionate amount of time on activities that aren't relationship-building: scanning industry news, scraping company information, maintaining target lists, building outreach sequences, and updating Customer Relationship Management (CRM) data. AI collapses the time spent on all of these.

- **Target screening.** AI agents can scan databases, news feeds, public filings, and trade press to identify companies matching your thesis. Per-target screening time drops from roughly 45 minutes to 8 minutes, freeing the team to evaluate dozens more opportunities per week.
- **Market mapping.** Build comprehensive sector maps with company profiles, ownership structures, growth metrics, and competitive positioning aggregated automatically from public sources.
- **Buyer and investor list construction.** Generate first-pass buyer lists by sector, deal size, geography, and recent transaction activity. The AI builds the long list. Your team curates it.
- **Outreach drafting.** Personalized cold outreach, follow-ups, and re-engagement sequences grounded in actual research about each contact and company rather than generic templates.
- **Relationship tracking.** Automate CRM hygiene: log meetings, summarize calls, identify relationship gaps, and surface prospects that have gone cold.

Pitch Development and Marketing Materials

Pitch creation is one of the most time-intensive activities in investment banking. It's also one of the highest-impact areas for AI, with Goldman Sachs reporting a 50% reduction in pitch creation time at scale.

- **Deck drafting.** Generate first-pass pitch books, management presentations, and teasers from a structured brief. AI produces an 80%-complete draft that your bankers refine rather than write from scratch.
- **Company and industry profiles.** Build the standard slide sets (company snapshot, business overview, transaction history, peer landscape, recent news) automatically from public sources.
- **Comparable company and precedent transaction tables.** Pull comp tables and precedents from public sources, calculate multiples, and flag outliers for review. The data still needs a banker's eyes, but the assembly work shrinks dramatically.
- **Tailored teasers and process letters.** Generate variants tailored to specific buyer types, sectors, or stages.
- **FINRA Rule 2210 review.** AI-generated retail and institutional communications still require supervisory review and approval, but AI can run a first-pass check against your firm's compliant language standards.

Due Diligence and Research

Due diligence is where time pressure meets information overload. AI doesn't replace the judgment calls. It handles the research, summarization, and document analysis that consume most of the hours.

- **CIM and document analysis.** Process a CIM in 30 to 60 minutes instead of 4 to 8 hours. AI extracts financials, customer concentrations, organizational structure, and identifies inconsistencies between the CIM and supporting documents.
- **Data room synthesis.** Summarize contracts, leases, customer agreements, and supplier agreements. Flag change-of-control provisions, exclusivity clauses, and unusual terms for legal review.
- **Earnings call and management interview analysis.** Transcribe, summarize, and extract themes from management calls. Compare against prior quarters or against peer commentary.
- **Industry and sector research.** Deep research across regulatory environments, competitive landscape, end-market trends, and capital structure benchmarks, with cited sources.
- **Financial statement analysis.** Upload a target P&L and get expert-level variance analysis, trend identification, anomaly detection, and scenario modeling.
- **Red-flag and red-team analysis.** Generate a structured list of reasons to be skeptical of the deal: data inconsistencies, customer concentration, legal exposures, market timing concerns.

Deal Execution

Once a process is live, AI helps move documents, tracks open items, and prepares the materials that consume the most analyst and associate hours.

- **Process management.** Maintain bid-tracking spreadsheets, working group lists, and status updates automatically. Generate weekly process updates for clients from the underlying data.
- **Document drafting.** First drafts of engagement letters, Indications of Interest (IOIs), letters of intent (LOIs), and process letters generated from a structured brief and prior templates.
- **Q&A management.** Triage and categorize buyer questions, draft preliminary responses for banker review, and track open items.
- **Closing checklists.** Generate transaction-specific closing checklists from past comparable deals. Update status automatically as items are completed.
- **Disclosure schedule support.** Initial review of disclosure schedule items against the purchase agreement for completeness and consistency.

Compliance, Control Room, and Supervision

Compliance is one of the highest-leverage areas for AI in a broker-dealer, both because the work is rules-driven and because regulators expect documented supervision. AI does not replace your Chief Compliance Officer's judgment, but it can dramatically reduce the manual work and improve consistency.

- **Marketing and communications review.** First-pass review of pitch materials, social media, and client communications against FINRA Rule 2210 standards. AI flags issues; humans approve.
- **Watch list and restricted list monitoring.** Cross-check client and deal team activity against firm watch lists and restricted lists.
- **MNPI and information barrier reminders.** Embed structured prompts and approval gates in workflows to remind staff before MNPI or confidential client data is entered into any system.
- **Policy maintenance.** Draft, update, and version AI use policies, cybersecurity policies, and supervisory procedures aligned to FINRA and SEC requirements.
- **Regulatory monitoring.** Track regulatory developments, FINRA notices, enforcement actions, and rule changes relevant to your registration and activities.
- **Books and records support.** Generate compliant audit trails for AI-assisted communications that may fall under SEC Rule 17a-4.

Practice Operations, Finance, and Billing

Back office work is high-volume, repetitive, and the kind of work AI excels at. The goal is to free your operations team for the judgment-driven work that matters.

- **Engagement letter and fee tracking.** Track retainers, success fees, expense reimbursements, and milestone billing across active engagements.
- **Time and project management.** Generate weekly utilization reports, project status updates, and capacity forecasts.
- **Financial reporting.** Build P&L reports, variance analyses, and revenue forecasts using historical engagement data.
- **Vendor and expense management.** Automate vendor due diligence reviews, expense categorization, and renewal tracking.
- **Document and email management.** Organize, tag, and retrieve documents across deals. Triage inbound email by deal, priority, and required action.

Human Resources and Talent

Hiring, retaining, and developing junior bankers is one of the most pressing strategic challenges in the industry. Attrition runs about 85% within 2 years at major firms. AI can't fix culture or compensation, but it can absorb much of the hours-intensive work that drives burnout.

- **Recruiting.** Draft job descriptions, screen resumes, build candidate comparison matrices, and prepare structured interview plans.
- **Onboarding.** Generate role-specific onboarding plans, training curricula, and 30/60/90 day check-in frameworks.
- **Training and development.** Build training content, prepare quick reference guides, and create coaching frameworks for managing directors and senior bankers.
- **Performance management.** Structure performance reviews, calibration prep, and development planning aligned to firm career frameworks.
- **Internal communications.** Draft firmwide announcements, deal team updates, and policy reminders that strike the right tone for an investment bank.

Measuring AI ROI

When evaluating AI investment, measure across 4 categories:

- **Time savings.** Hours recovered per banker per week. Track baseline (what each task used to take) versus current state. Use blended hourly cost to translate hours into dollars.
- **Error reduction.** Fewer typos in pitch books, fewer data inconsistencies, fewer missed items in disclosure schedules, fewer marketing pieces kicked back by compliance. Quantify where possible.
- **Revenue acceleration.** Faster pitch turnaround, more pitches per banker, more targets screened per week, faster deal execution. Track conversion rates and deal velocity.
- **Cost avoidance.** Lower spend on outsourced research, fewer compliance fines, lower turnover (especially among junior bankers), lower attrition replacement costs.

The 4-category framework matters because not every benefit shows up as a line item. The biggest gains often come from work that wouldn't have happened at all (the extra 20 targets you screened, the additional pitch you ran) because the team had capacity.

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 first-year analyst a blank deck and handing them a detailed pitch template with the standard slide architecture, the firm's brand guidelines, the comp table format, and a checklist of the items that must be in every book. Both can get the job done. 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 pitch builder agent, for example, doesn't just "draft a pitch." It works through a structured brief (sector, deal type, audience, key messages), pulls the standard slide framework, generates company and industry profiles in your house style, builds the initial comp set, and produces a draft with all the standard sections wired in for the banker to refine.

Guardrails. Agents include built-in quality controls. A target screening agent will cite its sources, flag low-confidence matches, and explicitly note when public information is too thin to draw conclusions. A compliance check agent will only flag, never approve, and route every output to a qualified supervisor.

Consistency. When 10 different bankers use the same agent, they get consistently structured outputs. The quality of a junior banker's first-pass deck doesn't depend on which managing director they were trained under.

Integration. Agents can connect to your business systems. They can read from your CRM, write to your deal pipeline tools, pull data from your document 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 for an investment banking firm:

- A workflow that monitors a target list for trigger events (earnings releases, leadership changes, financing activity, M&A news) and notifies the responsible banker with a structured summary and suggested outreach.
- A workflow that ingests an inbound CIM, generates a 1-page summary with key financials and deal terms, runs preliminary buyer fit scoring, and routes the package to the right coverage banker within minutes of receipt.
- A workflow that prepares the standard Monday-morning pipeline review automatically: deal status, hours by engagement, open issues, upcoming milestones, and red-flag items pulled from each banker's notes.

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. For bankers without prior automation experience, agents inside the platform can help architect and build new workflows through guided conversations.

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. They're easily cloned and customized for your firm's processes, terminology, and brand voice.

The library maps directly to investment banking functional areas:

- **Origination and deal sourcing:** Prospect research agent, deep research assistant, scenario planning strategist, brainstorming partner for thesis development.
- **Pitch and business development:** Presentation content creator, presentation slide generator, marketing image architect, B2B marketing letter agent, RFP builder, sales playbook curator, blog and thought leadership content creation.
- **Due diligence and research:** Deep research assistant with two-phase workflow and source verification, P&L financial analyst with Wall Street-level rigor, FP&A forecasting assistant, vendor assessment agent for diligence workstreams.
- **Compliance and supervision:** IT policy creator and editor, business continuity and disaster recovery plan agent, vendor assessment agent, knowledge base architect for policies and procedures, FAQ guru for the firm's compliance manual.

- **Practice operations:** Project plan architect, SOP writer, root cause analysis expert, change management advisor, EOS Level 10 meeting assistant for partner meetings.
- **HR and talent:** Job description specialist, hiring analyst pro, success profile builder, professional development plan agent, performance improvement plan agent, career coach.
- **Finance and back office:** P&L financial analyst, FP&A forecasting assistant, AI ROI break-even calculator.
- **Productivity:** Microsoft Excel expert, Microsoft Word support, prompt coach, devil's advocate, SMART goal setter.

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 agents tailored to your firm's specific deal processes, terminology, and house style.

CHAPTER 06

The Major AI Models and Providers

One Platform, Dozens of 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 have "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 dozens of 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 is added to the platform automatically.

The Major Providers

OpenAI builds the GPT family of models. These are the most widely recognized AI models and excel at general-purpose tasks including writing, analysis, coding, and creative work. GPT's smaller variants offer cost-efficient options for high-volume, simpler tasks.

Anthropic builds the Claude family. 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, which makes them well-suited for due diligence work and CIM analysis.

Google builds the Gemini family. Google's models bring strong multi-modal capabilities (handling text, images, and code together) and tight integration with Google's search and data infrastructure.

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

xAI builds the Grok family, high-capability models with strong reasoning and conversational abilities.

Amazon offers the Nova family, 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 complex multi-step analysis, mathematical work, and structured logic problems. These are especially useful for financial modeling reviews and quantitative analysis.

How to Choose the Right Model

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

- **High-stakes work** (client-facing deck drafts, fairness opinion analytics, complex financial reasoning, compliance reviews): Use frontier models from OpenAI, Anthropic, or Google. They cost more per query but produce the highest quality output.
- **High-volume work** (CRM hygiene, comp table assembly, news summarization, routine drafting): Use value models from Amazon, Meta, or the faster tiers from Google and OpenAI. They are significantly cheaper and fast.
- **Image generation:** The platform includes Nano Banana for creating custom images from text descriptions, useful for pitch book visuals, marketing materials, and conference graphics.

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

In a fast-moving market, the model that's leading today is rarely the model leading 6 months from now. The ability to switch between models with a single click (no new subscriptions, no rebuilt workflows, no migration projects) is one of the most underrated advantages of a multi-model platform.

For an investment bank specifically, this matters because different parts of the deal workflow favor different models. Drafting a CIM summary benefits from a model with strong reasoning over long documents. Building a comp set benefits from a model that's accurate with structured data. Generating a pitch headline benefits from a model with creative range. Lock yourself into one provider, and you're forcing every task through one set of strengths.

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:** Fast, cost-efficient image generation with real-time web search knowledge, precision text rendering, and translation for marketing mockups.
- **Nano Banana Pro:** Premium quality output with high-resolution options (up to 4K) and legible stylized text generation for infographics, conference materials, and pitch book visuals.

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. Like any communication skill, a few core principles go a long way.

The Fundamentals

Be specific about what you want. Instead of "Write me an email to a CEO," try "Write a 3-paragraph follow-up email to the CEO of a \$150M revenue industrial distribution company we met with at a sector conference. The tone should be professional but conversational. Reference a specific theme from the conference panel they participated on. End with a clear call to action to schedule a 30-minute follow-up next week."

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 managing director with 20 years of experience in middle-market industrial M&A" produces dramatically different output than a bare request to "analyze this company."

Provide examples. If you want output in a specific format or style, show the AI what good looks like. Paste in a prior CIM summary your firm produced, a previous pitch you liked, or an executive summary that matches your house voice. AI is excellent at pattern matching.

Break complex tasks into steps. Instead of asking AI to "build a pitch book," walk it through the process: "First, generate the company overview section using the brief I'm providing. Then build the industry landscape slides. Then draft the valuation framework discussion. Then assemble the comparable transactions slides." 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: "Tighten the tone." "Add a section on regulatory environment." "Shorten the executive summary to one page." 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 valuation range." This produces more thoughtful, accurate outputs and makes it easier to spot errors.

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 the deal team.

Constraint-based prompting. Tell AI what not to do. "Do not include forward-looking projections. Do not exceed 250 words. Do not include any company name that is not in the source document." Constraints sharpen the output and prevent common problems, including the kind that get flagged in Rule 2210 review.

Template prompts. Build reusable prompt templates for tasks your team performs regularly. A standardized prompt for drafting client status emails, building target screening summaries, or generating sector primers ensures consistent quality across the team.

Hallucinations and AI Accuracy

AI models occasionally produce content that sounds confident but is factually wrong. This is called a hallucination, and it's one of the most important risks to understand in a regulated industry.

For an investment bank, the consequences are amplified. A made-up precedent transaction in a pitch. A fabricated quote in a CIM summary. An invented regulatory citation in a memo. These are not just embarrassing. Under FINRA Rule 2210, communications must be fair, balanced, and not misleading whether they were generated by a human or a machine. Inaccurate AI output that reaches a client or a regulator is your firm's exposure, not the model's.

Reducing the Risk

- **Always verify facts, figures, and citations.** Treat AI output the way you would treat a junior banker's first draft: useful starting point, must be checked.
- **Use AI inside grounded workflows.** When AI agents are connected to authoritative sources (your house comp database, public filings databases, your firm's prior deal documents), hallucination rates drop sharply.
- **Prefer reasoning models for high-stakes work.** Frontier reasoning models are dramatically more accurate than legacy models.
- **Require source citations.** Build prompts that require AI to cite the source for every factual claim. If it can't cite, treat the claim as suspect.

Evaluating AI Output Quality

Before any AI output moves forward, ask 5 questions:

- **Is it accurate?** Are the facts, figures, and citations verifiable?
- **Is it complete?** Does it address everything the prompt asked for?
- **Is it on-brand?** Does it sound like your firm and your house style?
- **Is it compliant?** Does it meet FINRA Rule 2210 standards if it's a communication? Is anything misleading?
- **Is it usable?** Can a banker take this and finish it in a few minutes, or does it need to be rewritten?

The goal isn't 100% perfect output on the first prompt. The goal is an 80% draft that a banker can sand down to 100% in a fraction of the time it would have taken to start from blank.

Understanding Context Windows

Every AI model has a context window, the amount of text it can process in a single interaction. Modern frontier models can handle the equivalent of hundreds of pages, but performance can degrade as the window fills.

Practical tips for working in a banker's daily reality:

- For long documents (CIMs, 10-Ks, multi-deck pitch sets), summarize sections separately and synthesize the summaries.
- Start each new task in a fresh chat to prevent earlier context from contaminating the output.
- When asking for revisions, paste in only the relevant section, not the entire document.
- For analyses spanning multiple long documents, use an agent purpose-built for the workflow rather than chat.
- Keep an eye on the conversation length; if responses start to drift, start over with a tighter prompt.

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 (productivity, speed, competitive advantage) means nothing if your firm leaks MNPI, breaches client confidentiality, or shows up on a FINRA examiner's findings list because no one could explain how the firm governs AI.

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

The Shadow AI Problem in Investment Banking

Shadow AI is already happening inside investment banks. Analysts using ChatGPT to draft pitch summaries. Associates running comp tables through Gemini. Coverage bankers feeding client emails into Claude to clean up the tone. Operations staff using free tools to summarize vendor contracts.

The risk has nothing to do with banker curiosity. The risk is what they're putting into those tools. When an analyst pastes a deal memo, a CIM, a buyer list, or fairness opinion analysis 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 an investment bank, the categories of protected data make the exposure especially severe:

- **Material nonpublic information (MNPI).** Deal terms, valuations, board materials, pre-announcement transaction details. The legal exposure includes insider trading liability and Section 204A obligations of the Investment Advisers Act (where applicable).
- **Client confidential information.** Engagement letter terms, financial models, projections, customer lists, intellectual property disclosed under NDA.
- **Personally identifiable information (PII).** Executive personal information, ownership records, beneficial owner data.
- **Regulatory work product.** Marketing communications subject to FINRA Rule 2210, books and records subject to SEC Rule 17a-4.

When your cyber insurance carrier, your prime broker, your largest client, or a FINRA examiner asks how your firm governs AI, you need a documented, auditable answer. Not a shrug. A policy, a platform, and a paper trail.

The Investment Banking Compliance Landscape

For an investment bank, AI governance sits on top of a layered regulatory framework. The dominant frameworks for an SEC-registered broker-dealer doing M&A and capital markets work:

FINRA Rule 3110 (Supervision)

Member firms must establish and maintain supervisory systems reasonably designed to achieve compliance with securities laws. FINRA Regulatory Notice 25-07 (April 2025) extends Rule 3110 supervisory duties to generative AI workflows. Firms must establish written supervisory procedures specifically for AI applications, designate qualified supervisors for AI workflows, and maintain records of AI decision-making processes. FINRA Regulatory Notice 24-09 further required policies addressing technology governance, model risk management, data privacy and integrity, and model reliability.

FINRA Rule 4511 and SEC Rule 17a-4 (Books and Records)

AI-generated content that is transmitted (by email, chat, or otherwise) may constitute a communication subject to recordkeeping requirements. Firms must capture, preserve, and produce these records on demand. The SEC has proposed amendments to Rules 17a-3 and 17a-4 that would require books and records related to AI-driven conflicts of interest.

FINRA Rule 2210 (Communications with the Public)

All retail and institutional communications must be fair, balanced, and not misleading, regardless of whether they were authored by a human or by AI. AI-generated retail communications still require supervisory review and approval before distribution.

MNPI Handling, FINRA Rules 2241/2242, and Information Barriers

Investment banks must maintain documented procedures preventing the flow of MNPI between deal teams, research, sales and trading, and other lines of business. Any AI tool that touches deal information must respect those barriers, or it becomes the leak.

Bank Secrecy Act (BSA) and AML Program Requirements

Broker-dealers must maintain risk-based AML programs covering Customer Identification Programs (CIP), Know Your Customer (KYC) procedures, suspicious activity monitoring, OFAC sanctions screening, and ongoing customer due diligence. In March 2026, FinCEN imposed an \$80 million penalty on a broker-dealer for AML program failures, the largest BSA penalty ever against a broker-dealer.

Engagement Letter and NDA Confidentiality Obligations

Every client engagement is governed by signed engagement letter terms and, on the other side of every deal, NDAs that govern how target information may be used. AI tools that retain prompts indefinitely, train on customer data, or share data with third parties violate these obligations even if no regulator is watching.

Compliance Mapping: Where Managed Framework AI Lines Up

Regulatory Obligation	How Managed Framework AI Supports It
FINRA Rule 3110 (Supervision)	Centralized administration, full audit trails of prompts and responses, role-based access controls, written supervisory procedures supported by platform records
FINRA Rule 4511 / SEC Rule 17a-4 (Books and Records)	Searchable logs of AI communications, retention controls, audit trail export for examiners
FINRA Rule 2210 (Communications with the Public)	Supervisory review workflows, version control on marketing materials, structured approval gates before distribution
MNPI Handling and Information Barriers	Data isolation per organization, role-based permissions enforcing deal team segregation, no data retention by underlying AI providers
BSA/AML, KYC/CIP, OFAC	Encrypted storage of customer data, audit logs to support suspicious activity review, integration with screening tools through workflows
Engagement Letter and NDA Confidentiality	Contractual zero-data-training guarantees with every AI model provider, AES-256 encryption at rest, TLS 1.2+ in transit, no cross-organization data commingling
Vendor Due Diligence (FINRA Notice 21-29, 24-09)	SOC 2 Type I, Type II, and SOC 3 certification; documented security architecture; available SOC reports under NDA

The 6 Pillars of an AI Governance Framework

1. Policy. A written AI use policy defining approved tools, prohibited data types, MNPI handling rules, and acceptable use cases. Without a policy, you can't enforce your position or defend it in a FINRA exam. The policy doesn't need to be a 50-page document. It needs to exist, be communicated, and be enforceable.

2. Access control. Role-based permissions ensuring the right people have access to the right AI capabilities. A first-year analyst shouldn't have the same AI access as the head of capital markets. A well-designed permission structure prevents confusion, limits risk, and ensures deal team segregation 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 should guarantee that your data is never used to train any public AI model, full stop. This is what makes AI usable for an MNPI-handling firm in the first place.

4. Auditability. Full, searchable logs of who used AI, what prompts were entered, and what was returned. This is required by FINRA Rule 3110, increasingly expected by cyber insurers, and absolutely essential for any kind of internal investigation if something goes wrong. If you can't produce an AI governance report when asked, you have a gap.

5. Training. Ongoing education so your policies are understood and followed, not quietly worked around by well-meaning bankers. A policy memo doesn't count. Real training means structured onboarding, recurring support sessions, and a culture where people know the rules and the reasons behind them.

6. Control Room and MNPI Governance. Investment banks operate under deal team segregation requirements that don't exist in most industries. Your AI governance must explicitly address how MNPI flows (or doesn't flow) through AI tools, how watch list and restricted list controls apply to AI-assisted research, and how information barriers are preserved when AI agents are connected to firmwide data sources. This is the pillar that makes the other 5 work for a broker-dealer specifically.

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 firm's data is logically isolated. Your data is never commingled with other clients' data. Administrators have full visibility into usage, workflows, credit consumption, and audit trails, exactly what your Chief Compliance Officer needs to substantiate Rule 3110 supervision.

Compliance. The platform supports compliance with GDPR, CCPA, HIPAA (with Business Associate Agreement), and SOX. DPAs are available upon request. The certifications satisfy most enterprise security questionnaires and provide evidence of due diligence for cyber insurers, regulators, and client diligence teams.

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 a FINRA examiner asks how your firm governs AI, the answer should be documented, defensible, and supported by platform-generated records. That's what this architecture is designed to give you.

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.

The Roadmap at a Glance

Phase	Timeline	Focus	Key Milestones
Crawl	Weeks 1-4	Governance and foundation	AI policy, WSPs updated for Notice 25-07, DPA executed, users onboarded, champions identified
Walk	Months 1-4	Pilot workflows and first wins	Pre-built apps deployed, first pilot completed, ROI documented, use case repository started
Run	Months 4-6+	Scaling and self-sufficiency	Custom integrations, power users program, broader deployments, continuous governance loop

Why Crawl-Walk-Run?

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

- **Reduces risk.** Resolve governance, supervision, and MNPI questions before AI reaches a broad audience, not after a FINRA exam exposes a gap.
- **Builds confidence.** Give bankers time to learn the basics before asking them to automate complex workflows. 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 that funds further investment.
- **Creates internal expertise.** Each phase develops champions, power users, and eventually a community of practice, so the firm is never solely dependent on external help.
- **Scales sustainably.** By the time you reach the Run phase, the governance, cultural habits, and technical skills are already in place to support firmwide AI initiatives.

Crawl-Walk-Run is about going smart, not slow. The fastest path to transformational AI outcomes runs through a foundation of trust, capability, and supervisory rigor.

Phase 1: Crawl - Building the Foundation

Timeline: Weeks 1-4

The Crawl phase is about laying the groundwork. Before any banker builds a workflow or automates a process, the firm needs to answer fundamental questions: Who is allowed to use AI, and under what rules? How will MNPI be protected? How will Rule 3110 supervision apply to AI activity?

This phase may feel administrative, but it is indispensable. Firms that skip these steps invariably circle back to them later, usually under pressure, after a regulatory finding or a client concern.

1. Establish an AI Acceptable Use Policy aligned to FINRA and SEC obligations. Every AI journey begins with governance. Before the platform is turned on for a single banker, create or update a formal AI Acceptable Use Policy covering data handling, MNPI restrictions, information barrier preservation, FINRA Rule 2210 review requirements, and the boundaries of acceptable AI use. The policy doesn't need to be perfect on day 1. It needs to exist and be communicated.

2. Update the firm's written supervisory procedures. FINRA Regulatory Notice 25-07 requires written supervisory procedures specifically for AI workflows. Update your supervisory manual to reflect AI use cases, designated supervisors, and record retention requirements. This is the documentation the next FINRA exam will ask for.

3. Define roles, permissions, and your user structure. Define who will use the AI platform and what they will be able to do. A typical investment bank role structure includes administrators, Workshop/builder users, general users, and chat-only users. Organize users by deal team, function, or coverage group, respecting any information barrier requirements.

4. Execute the Data Processing Agreement and confirm books-and-records readiness. Sign the DPA. Confirm that AI usage logs satisfy your Rule 4511 retention requirements. Make sure the platform's audit logs are integrated with your firm's records preservation system.

5. Onboard users and deploy general-purpose AI applications. Bring users onto the platform through a controlled, deliberate process. Deploy a curated set of general-purpose applications to all users: writing assistants, summarization utilities, brainstorming aids, research helpers. Give everyone something immediately useful.

6. Create a central AI resource hub. Build an internal site dedicated to AI resources: getting-started guides, links to the platform, your AI use policy, a place to log wins, and contact information for support and Compliance.

7. Identify and develop AI champions. Select a small group of bankers and operations staff who are enthusiastic and influential, and invest in their skills early. Champions should complete the AI Champion Certification to build credibility and depth.

8. Establish an AI-forward culture from leadership. Culture change doesn't happen by memo. The managing partner, the head of capital markets, and the Chief Compliance Officer must visibly commit to AI adoption by using the tools themselves.

9. Launch recurring support sessions. Even with great documentation, people need live support. This is where Framework IT's Office Hours come in. Recurring, low-pressure sessions where users can ask questions, troubleshoot prompts, and share discoveries.

By the end of Crawl: Governance framework in place, supervisory procedures updated, users onboarded and exploring, champions developing their expertise, and leadership visibly backing the effort.

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 is ready to tackle real workflows with guidance. This phase is defined by intentional, supported experimentation: selecting the right problems, building solutions collaboratively, testing in controlled environments, and measuring results.

1. Review and deploy pre-built applications. Strong starting points for an investment bank: SOP Writer (for supervisory procedures), IT Policy Creator (for AI policy), Deep Research Assistant (for sector primers and target screening), P&L Financial Analyst (for target analysis), Compliance Check (for first-pass Rule 2210 review).

2. Customize applications for specific teams. Introduce specialized applications tailored to specific functions. Schedule working sessions to walk through how each tool would be used in practice.

3. Identify your first pilot workflow. The ideal pilot is high value, low complexity, visible, and measurable. Strong pilot candidates: target screening workflow, CIM summary workflow, weekly pipeline review workflow, sector primer generation.

4. Develop the AI solution with guidance. Build the pilot solution as a collaborative, guided effort. The bankers and operations staff involved in building the pilot develop skills they carry into future projects.

5. Test with a controlled pilot group. Deploy the solution to a small group first: a single coverage team, a single deal team, or a handful of volunteers.

6. Measure ROI and document your first win. After the pilot has run long enough, measure results against your baseline. Share the results. Present them to leadership. This first documented win is the proof point that turns skeptics into supporters.

7. Maintain a living use case repository. Capture every AI use case, successful and unsuccessful, in a structured document.

By the end of Walk: Specialized AI applications deployed to targeted teams, first pilot automation completed and measured, ROI documented, and a repeatable process established for future use cases.

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

Timeline: Months 4-6+

The Run phase marks the transition from guided experimentation to firmwide self-sufficiency. The guardrails are still in place, but the firm is driving.

- 1. Develop a power users program.** Identify bankers and operations professionals who have demonstrated strong AI skills and a willingness to help others. Power users serve as a distributed support network.
- 2. Explore custom integrations.** Connect AI to your CRM, deal pipeline tools, document management, and communication platforms. These integrations unlock dramatically more powerful workflows.
- 3. Formalize business case analysis.** Move from informal estimates to formal business case analysis for new AI initiatives.
- 4. Enable proactive workflow discovery.** Teams should be identifying automation opportunities on their own.
- 5. Build custom AI applications.** Your team is now capable of building custom agents and workflows that go beyond the standard library to solve challenges unique to your firm.
- 6. Execute broader deployments.** Roll out successful solutions to additional groups, offices, or coverage areas.
- 7. Track ongoing value and adoption metrics.** Use your AI hub, platform analytics, and use case repository to continuously monitor adoption rates and value.
- 8. Establish a continuous improvement and governance loop.** Create a formal, recurring forum where the managing partner, Chief Compliance Officer, power users, and program managers review adoption metrics, prioritize new use cases, and address governance updates.

By the end of Run (which never truly ends): Your firm is building and deploying AI solutions independently, tracking measurable value across coverage teams and operations, and continuously identifying new opportunities. AI is no longer a project. It is part of how the firm works.

The People Side of AI Adoption

Technology adoption is 20% technology and 80% people. The firms that succeed treat AI rollout as a change management initiative, not a software implementation.

- **Visible leadership use.** Managing partners and Chief Compliance Officers who use AI themselves create permission and momentum throughout the firm. Bankers watch what leaders do, not what they say.
- **Honor the skeptics.** Senior bankers who built their careers without AI deserve a respectful conversation about how this fits into their work. Often their concerns identify gaps that strengthen the rollout.

- **Celebrate the early wins publicly.** When an associate cuts their pitch prep time in half, share it. The example travels.
- **Make compliance part of the celebration.** Show how the governed AI program is making the firm safer, not just faster. That message resonates with bankers and senior leadership alike.

Data Readiness: Garbage In, Garbage Out

AI outputs are only as good as the data they're fed. Before scaling AI use, audit:

- **Document quality.** Are your deal templates, comp databases, and house style guides up to date and accessible?
- **System integration.** Are your CRM, deal pipeline, and document storage systems organized in ways that AI tools can read from?
- **Source authority.** Are your bankers using authoritative internal sources, or are they reinventing the wheel for every pitch?

Clean data is the difference between AI that drives a 50% improvement and AI that produces output the team has to redo by hand.

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 investment banks adopt AI in a way that reduces risk around MNPI exposure, client confidentiality, supervision, and policy misalignment. Leadership can approve AI with confidence. Bankers can use it without creating invisible liability.

Governed. We give the Chief Compliance Officer and the managing partner visibility, control, and audit trails so AI use is managed intentionally rather than emerging chaotically. When a FINRA examiner, a cyber insurance carrier, or a board member asks "How is AI governed here?" your firm has a documented, defensible answer.

Productive. We turn AI into practical business outcomes through enablement, use-case alignment, workflow improvement, and measurable gains in deal velocity, pitch quality, and analyst capacity. The investment in AI compounds over time instead of quietly gathering dust.

What Is Included

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

Unlimited Users. No per-seat pricing surprises. Every banker, every operations professional, every member of the compliance team included.

100+ Framework IT Standard Agents and Workflows. Purpose-built tools covering origination, pitch, due diligence, deal execution, compliance, operations, finance, and HR. Easily cloned and customized for your firm. Ready on day 1, in addition to 500+ community-built apps on the platform.

No-Code Workflow Builder. Build powerful automations with drag-and-drop. 50+ 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.

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. Ongoing role-specific training resources that reinforce the certification.

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, and they are one of the most valuable parts of the entire program.

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 an analyst is struggling with a prompt, an associate is trying to figure out which model to use, or a Chief Compliance Officer is wondering how to approach Rule 2210 review of AI-generated content, our team works through it in real time. No question is too basic.

Use Case Workshops. Structured sessions where we work through building a specific workflow, agent, or automation together. We build live, with your input, on problems that matter to your firm. Bankers can bring their own use cases (a pitch template, a CIM screening workflow, a buyer list builder) and we'll help architect and build custom agents or workflows right there in the session.

Show and Tell and Peer Learning. Users share wins, discoveries, tips, and creative applications, facilitated by our team. The collective intelligence of the group accelerates everyone's learning.

Office Hours are low-pressure, inclusive, and welcoming to beginners. They serve a dual purpose: accelerating individual skill development and surfacing the use cases and pain points that inform your firm's AI roadmap.

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 and continuously expanding.

During each SBR, we:

- **Review adoption metrics.** Active users, usage patterns, credit consumption, and engagement trends.
- **Assess progress against your AI adoption roadmap.** Where are you in the Crawl-Walk-Run framework? What milestones have been hit?
- **Identify new use cases and automation opportunities.** Based on what your team has learned, what new workflows should be built?
- **Update and refine your AI roadmap.** Your roadmap is a living plan that evolves as your firm's AI maturity grows.
- **Plan training and enablement activities.** What Office Hours topics should we prioritize?

The SBR keeps AI adoption moving forward as an ongoing strategic initiative with a named partner who stays 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.

The key word is "with." Accelerator Plans are collaborative engagements designed so your team learns and upskills along the way. We build together so your people develop the expertise to keep building on their own.

Typical Accelerator Plan engagements include:

- **Custom agent development.** Pitch builder agents, CIM analysis agents, target screening agents, sector primer agents, fairness opinion support agents (analytical, not legal).
- **Workflow automation.** Multi-step automations that connect AI to your CRM, deal pipeline, and document storage, from scoping and architecture through testing and deployment.
- **Function-specific rollouts.** Working with a specific group (M&A advisory, capital markets, restructuring) to identify their highest-value AI use cases, build the solutions, train the team, and measure results.
- **Compliance integration.** Building supervisory review workflows, Rule 2210 first-pass screens, and audit trail exports tailored to your firm's existing compliance manual.

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
Pre-Built Agents	Limited	Limited	100+ Framework IT + 500+ community
FINRA Rule 3110 Audit Trail	Limited	No	Full supervisory audit trail
MNPI / Control Room Controls	No	No	Role-based access, data isolation, supervisory workflows

The Transformation

Before Managed Framework AI: The firm is operating in the dark. Analysts and associates are using AI tools nobody approved. The Chief Compliance Officer has no visibility into what is being exposed. The managing partner can't answer the question of how AI is governed. Pitch quality is inconsistent. Rule 2210 review is harder than it should be. There's skepticism, risk, and quiet anxiety.

After Managed Framework AI: The firm has a named partner, a documented roadmap, and a trained team actively using governed AI every day. Outputs are consistent and trusted. Pitch creation time is dropping. CIM analysis is faster. Compliance has a documented answer for every regulator question. Leadership can see adoption dashboards and prove ROI. 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 you stand 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 managing partner and Chief Compliance Officer have formally endorsed our AI strategy
- We have designated an internal AI Champion to lead adoption
- Our AI policy has been communicated to all bankers and operations staff
- All employees have acknowledged receipt and understanding of the AI policy
- Our AI policy is reviewed and updated at least annually
- Senior leadership actively uses AI tools and visibly supports the initiative

FINRA, SEC, and Broker-Dealer Compliance

- Our written supervisory procedures have been updated to address AI workflows per Regulatory Notice 25-07
- We have designated qualified supervisors for AI activities under FINRA Rule 3110
- AI-generated communications subject to FINRA Rule 2210 go through documented supervisory review
- We maintain books and records of AI-related communications per Rule 4511 and SEC Rule 17a-4
- Our AI use is included in our annual compliance certification and risk assessment

- We have documented MNPI handling rules for AI tools and information barriers
- Our AI vendor due diligence aligns with FINRA Notice 21-29 and 24-09 requirements

AML, KYC, and Confidentiality

- Our AML, KYC/CIP, and OFAC programs explicitly address AI tool usage
- Client engagement letters and NDAs are reviewed against AI tool capabilities
- We have documented restrictions on entering MNPI, client confidential information, and PII into AI tools
- Our cyber insurance carrier is aware of our AI usage and governance framework
- We have a Data Processing Agreement in place with our AI platform provider

Access Control and Permissions

- AI access is role-based, not everyone has the same permissions
- Deal team segregation and information barriers are enforced in AI tools
- We maintain searchable logs of AI usage (who, what, when)
- We can produce an AI governance report on demand for auditors or examiners
- We have a process for reviewing and approving new AI tools before deployment
- User permissions are reviewed and updated when roles change
- 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 across business lines
- 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 all client-facing and supervised 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 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 supervision.

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 groups, 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 will walk through where your firm stands today, identify the highest-value opportunities (typically pitch automation, CIM analysis, and target screening for new clients), and show you exactly how Managed Framework AI works. We can also work through how the platform supports your FINRA Rule 3110 supervisory obligations and your firm's MNPI handling requirements. No pressure, no pitch deck. Just a conversation about what makes sense for your firm.

Take the readiness checklist to your leadership team. Use it to start an internal conversation about AI governance, supervision, and opportunity. The checklist alone is often enough to surface gaps that nobody was talking about. Bring it to your next compliance committee or management committee meeting.

Stop the bleeding on Shadow AI. Every day your analysts and associates use unmanaged AI tools is another day MNPI, client information, and pitch material are 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 are here to help you write it.

Book a Consultation

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