Abacum

Guide

AI: Faster, Smarter Decisions.

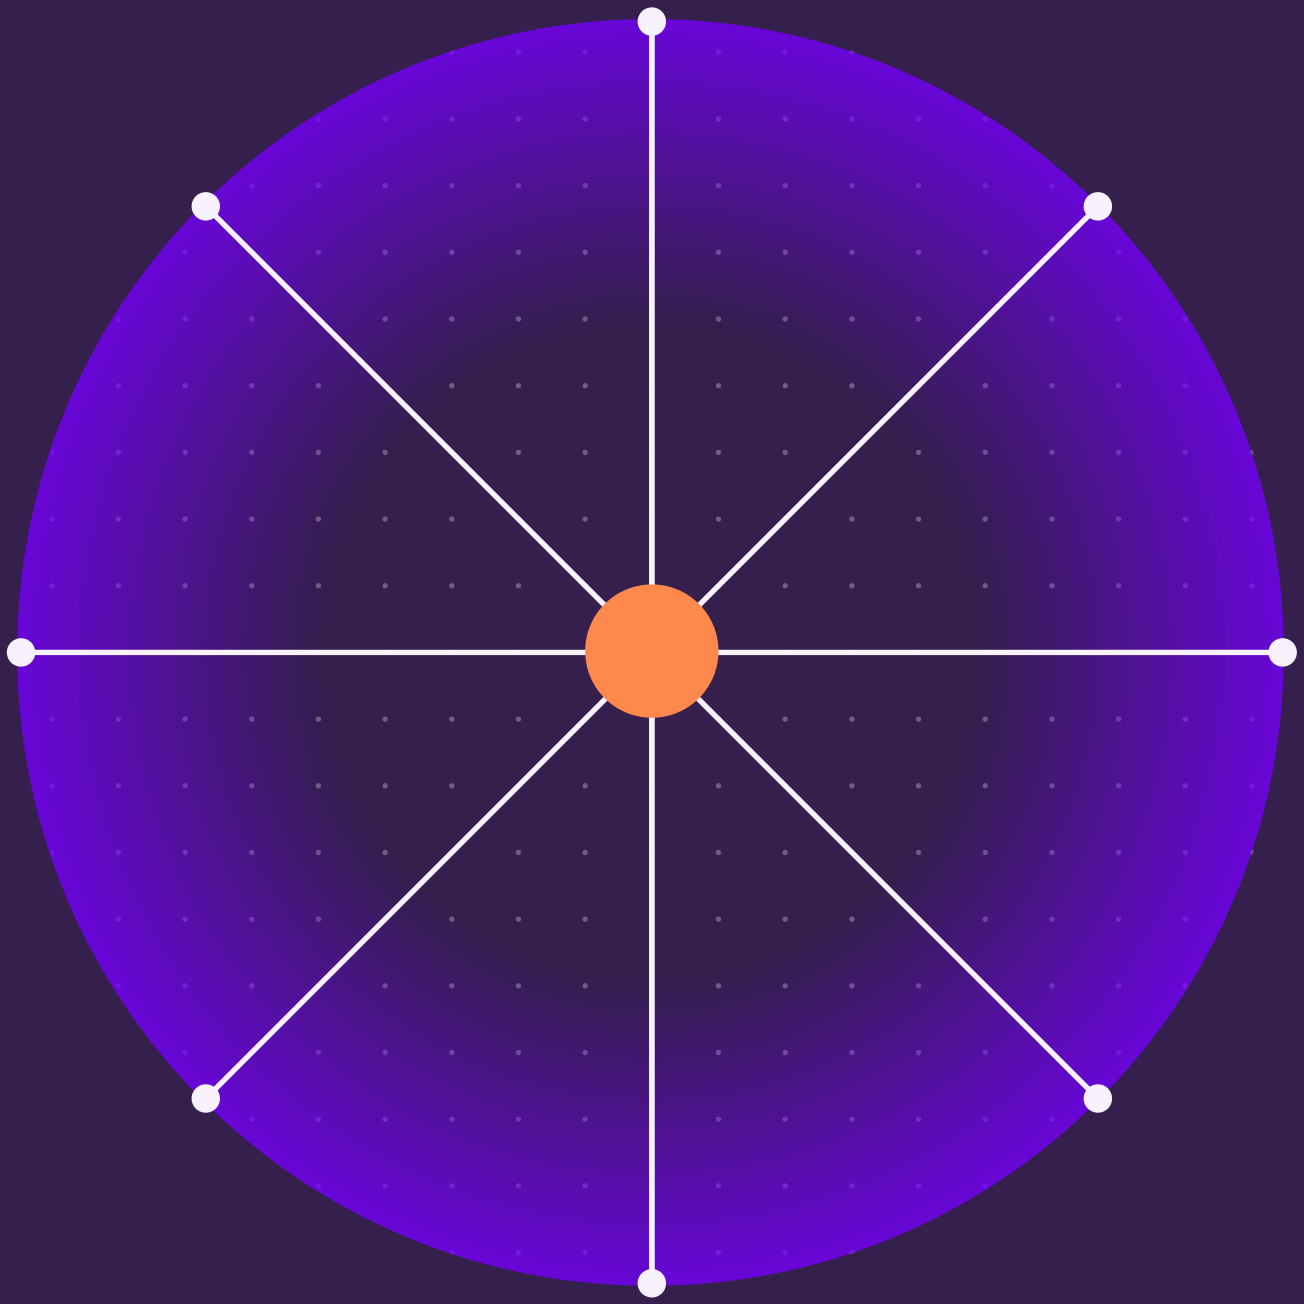

Table of contents

Foreword

This guide is not for everyone. We have written it for finance professionals who are dealing with the day-to-day chaos of leading rapidly scaling mid-sized companies. We know it's messy.

This is our second guide. Now, we get into the specifics of actually using AI. We begin with **1. What to expect from AI today** to give an overview of what is possible. Then we get into actually **2. How to roll out "AI," properly** before discussing **3. Monitoring Performance and ROI**. We then end with discussions on **4. Why Finance needs to take charge** and **5. AI regulations**, fun times.

The best teams will use AI as it is meant to be used, a strong tool that can improve the speed of making decisions when properly used. To do so, this means centralizing your data and implementing in a quick centralized fashion. We hope to save you time.

We have read all the books. We know the best practices. But what makes our guides unique is our experience of being in the day-to-day finance trenches at scale-ups and start-ups across the world. At Abacum, we've drawn from hundreds of real-world missteps to provide grounded advice that enables CFOs to drive effective growth. We've minimized the fluff, added in first-person anecdotal context, and focused on tactics you can directly implement. We are finance professionals building for other finance professionals

Who we are

Abacum was founded by Julio Martinez and Jorge Lluch, former CEO/CFOs who set out to create the solution they always wished they had. They recognized a critical problem: the FP&A process is fundamentally broken. Finance teams spend too much time on tedious manual tasks, leaving little room to focus on strategic decision-making and driving faster execution.

To solve this, they put together knowledge from hundreds of FP&A specialists with a vast range of expertise, from Fortune 500s to pre-seed successes, to build Abacum, an FP&A platform to drive business performance.

Our stories and advice are based on insights from a variety of discussions, experiences, and observations from our own professional backgrounds and from helping hundreds of finance teams improve their FP&A processes. They are not specific to any particular company.

1. What to expect from AI today

The FP&A AI dream...

AI is a tool. And like every tool before it, it will solve some of your problems in a novel way. Take double-entry accounting, invented in Italy in 1494. It allowed businesses to thrive by building trust in financial systems. Later came Excel in 1985, making analysis and financial modeling accessible. In 1990, ERPs transformed how companies operated by integrating every major business function. Most companies have been using parts of AI for years (excluding invoicing automation which has been around since 1965).

But here's the truth:

> None of these tools make strategic decisions, and AI shouldn't either. At the end of the day, it will always be people who make decisions about directions and priorities. And, It is Finance's job to enable that by providing strong data.

When finance embraced double-entry accounting, ERPs, and Excel, the function grew into the strategic powerhouse it is today. And the teams that embraced these tools early moved faster and achieved success more quickly.

So, paradoxically, AI is both the most underhyped and overhyped technology of our time. Often, this confusion leads to severely mis-matched expectations.

We start our first section in our second guide, Deploying AI for FP&A, with answering this question by a framework. Then, we will walk through each major task of an FP&A team to discuss the ideal goal for AI use, ways to use AI now, and then a best in class example.

In terms of a framework, AI's general capabilities are:

Great	Automating tasks with clear rules (e.g., data cleanup, consolidation, accounting entries)
Good	Giving quick answers to questions where data is structured (e.g., analyses, reports, scenario planning, alert systems) thanks to chat-based interfaces
Not Great	Tackling problems with high uncertainty or requiring persuasion
Awful	Making strategic decisions

And now, how this applies across FP&A work:

FP&A Time	Wins	How	Time saved	Ideal examples
1. Data Consolidation & Reconciliation	Auto-cleaned imports, error detection, and data alignment.	Auto-import key systems, flag issues, guide fixes, and monitor ongoingly.	Eliminates reconciliation, ensures system trust and historical consistency.	Full system sync with AI-driven checks and alerts (e.g., $20k ARR drop flagged and resolved).
2. Reporting	Auto-generated reports, narrative insights, multi-version output, and historical context.	Describe and auto-create reports, explain insights, share interactively.	Simplifies multi-report reviews, ensures alignment, and automates sharing.	Auto-updating, audience-tailored reports with consistent narratives (e.g., board, EOM, weekly).
3. Planning	Automated models, intuitive budget walkthroughs.	Define goals, build baseline, iterate with co-pilot, and cascade actions.	Speeds up model building, ensures accuracy, and centralizes updates.	Rapidly built models with action-based guidance (e.g., budget, runway, metric insights).
4. Scenarios	Fast, with decision tracking and follow-up what-if questions	Pose questions, simulate, propagate changes, and track decisions.	Handles complex scenario impacts quickly and cohesively.	Live decision tracking with impact modeling (e.g., churn increase scenario).

FP&A Time	Wins	How	Time saved	Ideal examples
5. **Agents / Business Partners**	Real-time, tailored insights for teams and leaders.	Daily summaries, instant alerts, and proactive issue flagging.	Cuts down daily reporting burdens and unifies leadership communication.	Morning briefings with deep insight, alerts, and proactive next steps (e.g., pipeline anomaly fix).
6. **Memory (Version History)**	Instant access to historical decisions, tracked outcomes, and narrative continuity.	Log decisions, re-run reports with past logic, generate performance narratives, and store insights.	Removes the need for manual retrospectives and re-analysis; decisions and results are instantly retrievable.	AI retrieves past pricing decisions, compares outcomes, and provides narratives on what worked and what didn't.

Important:

1. Nearly all of AI's value in FP&A derives from having all of the information from all your systems in one place, which is critical for nearly any FP&A work.
2. Your current data stack will define how you use AI. If it's in Excel, load it into an LLM and then generate a Python file if you want to use it. If it's across large Google Sheets, connect a big LLM model to your sheets through an API. If it's in a data warehouse, integrate tools over the platform. If it's in an FP&A tool, expect most of these features are either present or on the roadmap.

Lead with AI's strengths: automating repetitive actions and allowing chat interfaces to provide truer understanding

Layer on more and more automations as your centralized data becomes cleaner.

The end goal should be saving you time and getting to solutions more quickly, not raising more questions.

Let's get to work.

1 Data Consolidation & Consolidation

People have been harping on about data consolidation since the first ERP. Then, consultants began selling the wonderful benefits of integrations. But the sad fact is, forty years later, finance teams still spend a ridiculous amount of time reconciling systems, importing spreadsheets, and fixing data errors.

Goal: All your systems are the source of truth. AI constantly checks across them all to give you perfection.

Ways AI can help now:

Cleaning data for Import	For a single import, you can use simple tools to sift through thousands of transactions, identify errors, and suggest fixes across widely different formats. A good system will use AI to validate across all of your systems and uncover underlying issues. **Example:** Import all of your CRM, ERP, HRIS, and billing data. Define issues: Customer X is missing in HubSpot, Employee Y isn't reflected in the HRIS from last June, or 15 companies show no usage but haven't paid bills and have technically churned.
Monitoring issues	AI agents can run on top of each system and across them as an overview layer. They can tell you when, say, a customer has the wrong ID in your billing database.
Fixing issues	AI agents can propose and execute fixes, once approved, by matching across several criteria to clean transactions quickly and accurately.

You know you're doing it well when this is automated:

> "The CRM logs a $20k ARR drop due to a churn. The system checks billing and ERP to confirm alignment. It checks whether this is within expected churn. If not, it flags a Slack channel with recommended actions, then confirms once the issue is fixed."

Done well, AI checking your data and proposing fixes is like having someone reconciling all your data non-stop. Once the data is in a good place, we can turn to reporting.

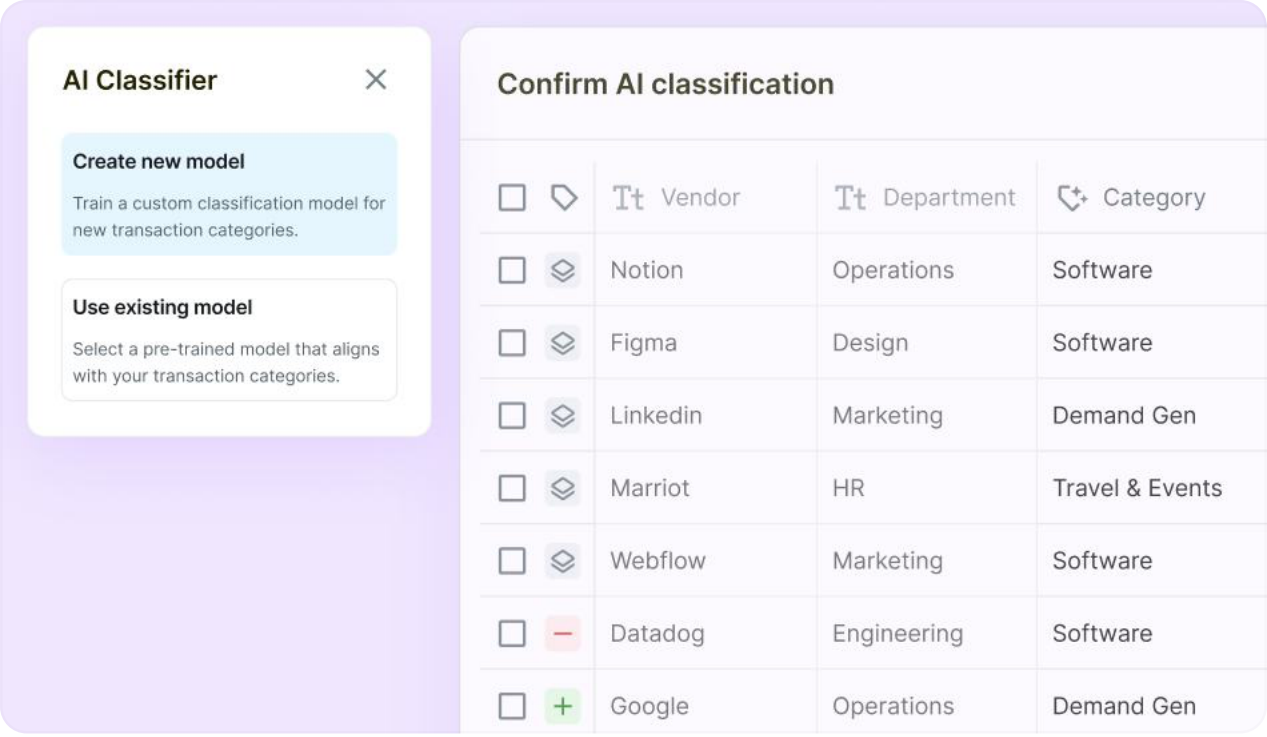

If you're hesitant, try this: upload a particularly tough reconciliation to a general LLM (Google, OpenAI, etc.) and ask it to analyze.

2 Reporting

Reporting is a balance between showing a single version of the truth (the company view) and slicing reports in different ways to answer specific questions. Traditional reports always face this tension, sometimes solved with drilldowns and filters, which inevitably leads to people pulling the wrong data.

Goal: All slices of reports and decks match effortlessly, past, present, and future. Anyone can ask a question, like "How is my customer, Perplexify, doing in terms of usage and MRR over time?"

Ways AI can help now:

Chat-based user interface	People are increasingly comfortable asking questions in context. This UI is much friendlier than clicking through dashboards.
Generate reports	Creating one report isn't that time-consuming. But what about 20 slight variations? Or reformatting for the board? Or pushing the same metrics to 20 different stakeholders? AI can cut this time from weeks to hours.
Look beneath the reports	Perhaps the most interesting capability. Reports often struggle with answering "Why did sales increase by 10%?" AI can now help. (E.g., "X new customers bought Y product; marketing conversion rose Z%.") When done right, these systems spot the actual events without needing to drill-down for hours.
Automate explanations	Many tools now summarize reports to help those less comfortable with numbers understand what's going on.
Take KPIs to the next level	You can track more data, and track more relevant data without needing frequent updates, which actually allows people to use KPIs effectively.

You know you're doing it well when this is automated:

> "Sales conversion this past week is 10% lower than average, even adjusted for seasonality. Root cause: fewer ICP leads routed to the Europe sales team based on two underperforming marketing campaigns" Or an immediate answer to: "Break down margin by region, product, and month for the past two years."

Summarizing large volumes of data (and let's be honest, this was mostly done with SUMIFs) is something AI does incredibly well.

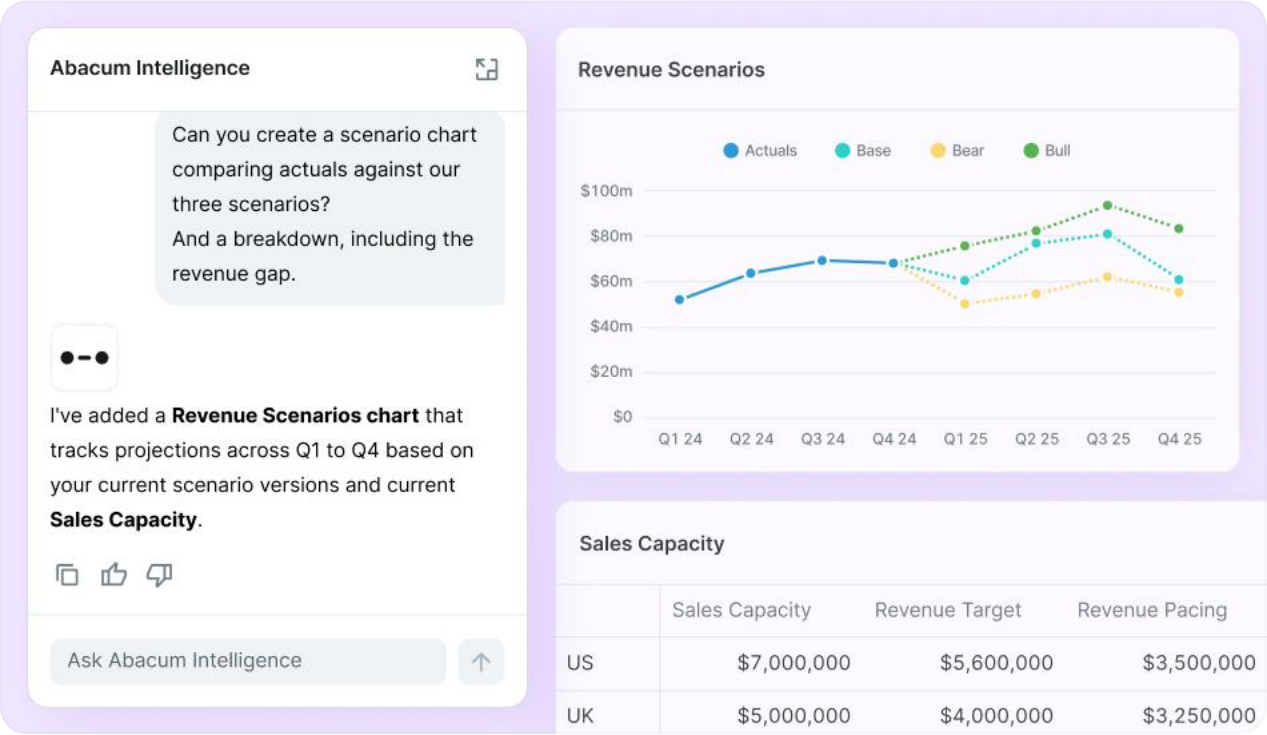

Building too many models upfront can miss the point. First, make sure there's a central repository where major models are validated, controlled, and created. Otherwise, you risk having your data tell completely different stories.

3 Scenarios

Building on reporting, you'll eventually want to explore future scenarios, whether it's budgeting, pricing, or go-to-market strategies. Like reporting, you can generate these scenarios easily with the right tools.

Goal: Answer any *What If?* question instantly. Track hundreds of decisions live to evaluate their impact.

Ways AI can help now:

Question, Answer	Pair question-asking with AI-generated reporting, and your CEO can get answers to the kinds of questions that drive real decisions.
Tweaking many variables	Generative AI can iterate thousands of scenarios across hundreds of variables. It helps you choose the most realistic paths—and shows which assumptions might derail you.
And much more	Planning and budgeting...

You know you're doing it well when this is automated:

> "Show me a downside scenario where next quarter's new sales are 20% below plan, churn increases by 2 percentage points, and we roll out the new import feature as an add-on to customers who joined in the last year."

By saving time on building scenarios, your team can focus on which scenario is most realistic and what targets are truly worth aiming for.

4 Agents/Business Partners

Okay, now you have the data consolidated, the reports set up, and you can run scenarios at a touch. It's time to let agents act as analysts across the business.

Goal: Everyone on the team stays aware of how the business is performing, with real-time alerts on what needs attention.

Ways AI can help now:

Monitor rules	Set up AI to monitor rules across systems and send alerts. For example, if reconciliations are off or cash from a key customer is late.
Work as business partners	Assign one agent per team to field data questions, send tailored reports, and provide a custom dashboard view.
Receivables partners	Similar to customer service bots, agents can send overdue payment reminders. These can be customized with language based on usage trends, making them more effective.
An extra example	BCG report on generative AI in the Finance Function

You know you're doing it well when this is automated:

> "Good morning! Here is your Morning Report:
> - Cash balance is $12.4M (down $0.3M day-over-day, reflecting yesterday's payroll run).
> - ARR: $50.2M, up $200k from new bookings yesterday (2 deals closed in Salesforce).
> - Churn: No cancellations recorded yesterday. Current MoM churn is 1.8%, slightly above our 1.5% target.

> - Expenses: Marketing spend month-to-date is $1.2M, which is 10% over the run-rate budget. Mainly due to higher “Dreaming of Integrations” campaign.
> - Forecast vs Plan: We are tracking 3% below plan on revenue for Q3. At this pace, end-of-quarter revenue is projected to be $X, which is $Y below target.
> - Today’s Insight: An anomaly was detected in the Europe sales pipeline – conversion rates dropped to 15% this week vs. 25% average. This might need sales leadership attention.”

Now, we aren’t there yet. We still need analysts. But, with current trends this is highly likely over the coming year. For now, the best way is to have an AI agent running this across your system with an analyst checking and then distributing to the team.

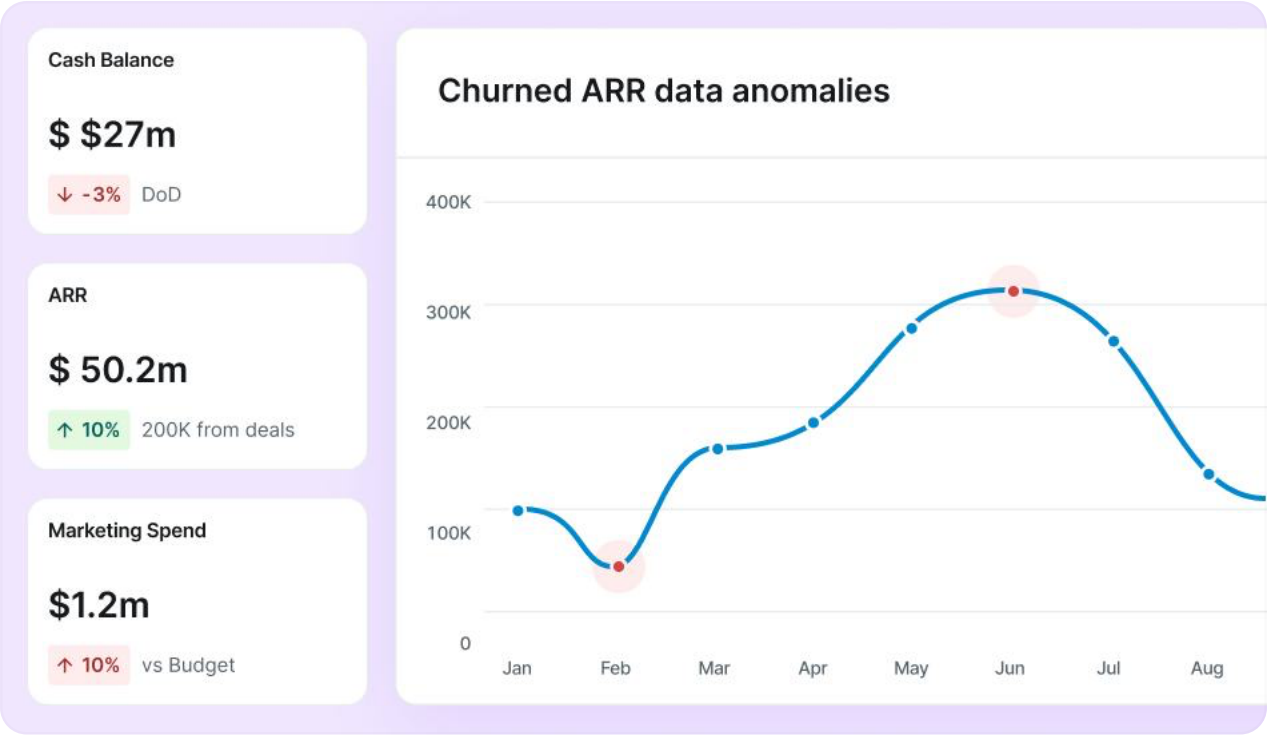

Ensure AI agents actually add value. The last thing you want is 50 *Clippies* running around your financial systems.

5 Memory

Over time, the right setup helps you tell a consistent story. You'll be able to quickly look at how reports appeared in the past, track the decisions made, and analyze the outcomes. This can be intensely time-consuming for individuals, but when systematized, that knowledge becomes instantly accessible.

Goal: All current decisions are informed by past decisions and their outcomes without wading through endless tables that no one reads.

Ways AI can help now:

Troubleshoot past reports	You can review historical reports and compare them with re-runs using the same logic. This is especially useful in due diligence.
Construct a narrative	Since the system understands past data, you can ask it questions like: "How has sales commission changed over the past five years? What results did it drive? How does it compare to today?"
Hold teams accountable	Agents can generate performance summaries quickly. For example, they could review the past six months of a sales rep's activity and produce a data-backed performance review.

You know you're doing it well when this is automated:

> Every decision you make has a clear, updated narrative including previous decisions.

In conclusion

AI is good at some things, not great at others. Lean into where it is genuinely useful for its ability to interpret words quickly, process thousands of transactions instantly, and interact through text with your team.

Yes, it takes some effort to set up. But like any good system implementation, once it's in place, the payoff is faster, smarter decisions.

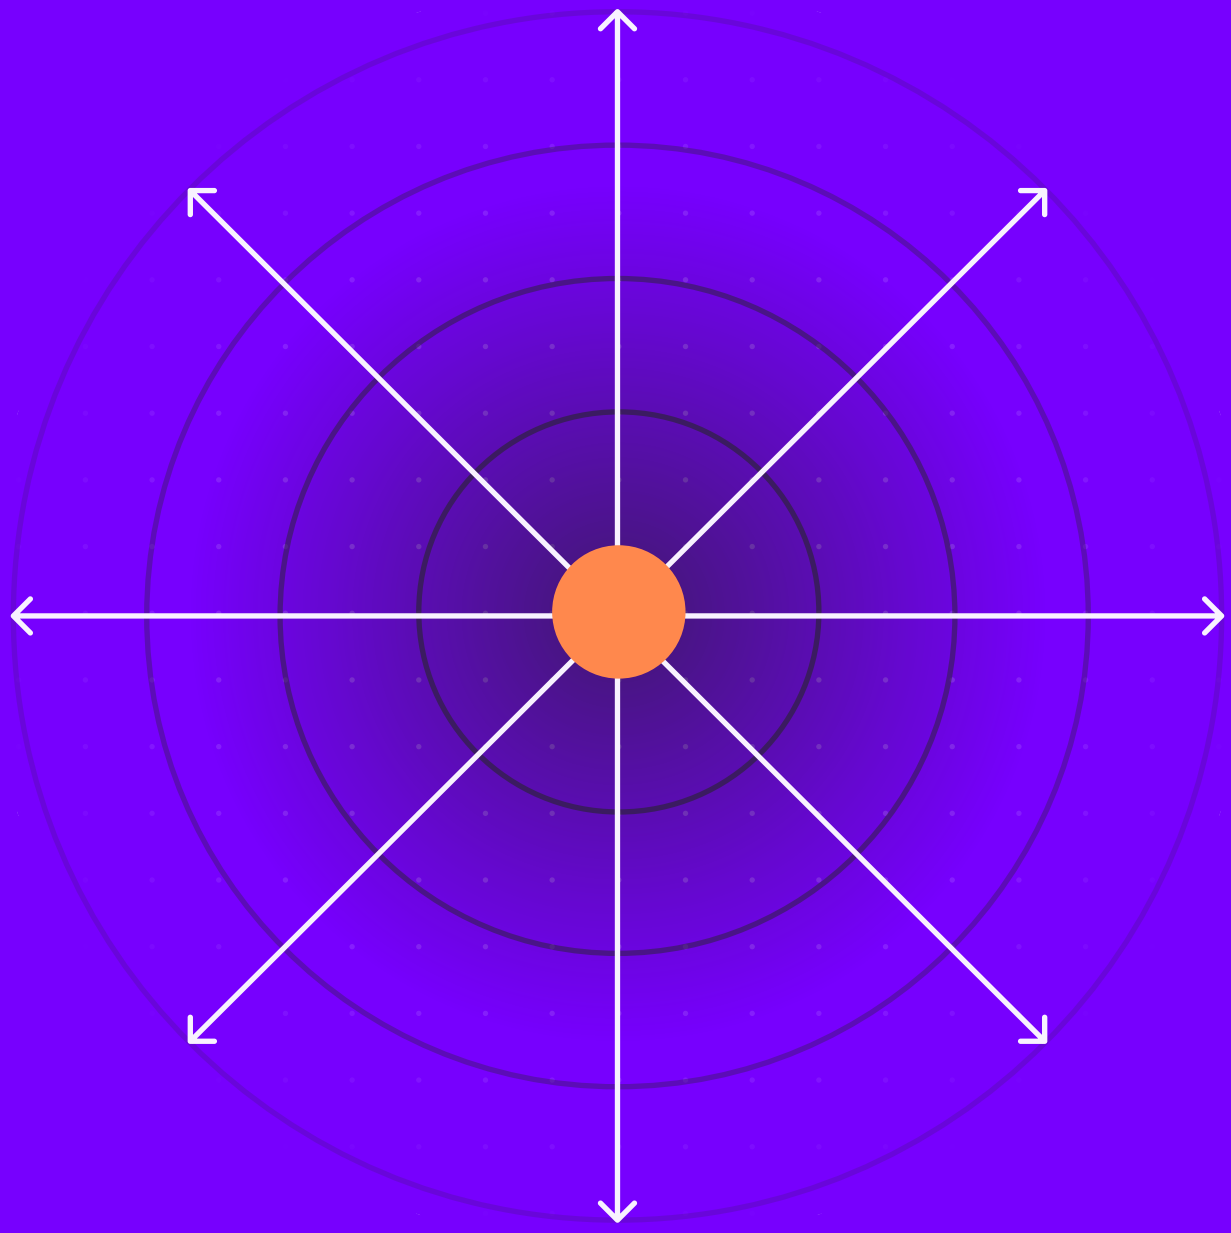

2. How to roll out “AI” properly

Good companies do AI right, and they are increasing their lead

In the last section, we discussed the dream of AI. But like any other tech, it takes the right setup to make it useful. History is littered with good intentioned finance systems, gone wrong:

There was Hershey Foods’ infamous 1999 ERP and CRM go-live disaster over Halloween resulting in over $100 million of undelivered inventory. Cadbury Schweppes’ 2006 ERP rollout led to a $12 million chocolate inventory pile-up from bad order projections. Mission Produce’s ERP launch? It resulted in a complete loss of visibility into how their avocados were ripening, costing millions in spoiled fruit and the indignity of having to buy the fruit from competitors.

But it’s not just chocolate and avocados.

In 2017, Amazon.com spent two years implementing a cloud-based HRIS system, which quietly failed. And then there was good ol’ JP Morgan’s “London Whale” incident? A simple formula error in an Excel-based VaR model, dividing by the sum of rates instead of the average, contributed to a $2 billion loss by drastically underestimating risk. (And some more fun ones).

> Currently, nearly 42% of big companies that have piloted AI projects have now abandoned them.

It’s essential to get this it right. When implemented properly, it can payback in 6–12 months. These benefits compound, driving dramatically better performance over time.

As we outlined in the last section, AI’s core strengths for FP&A right now are in two areas:

1. **Automating rule-based tasks**	Think data cleanup, consolidations, accounting entries.
2. **Answering questions quickly**	Analyses, reports, scenario planning, alerting in user-friendly chat-based interfaces that work across **a lot** of data.

So how do you solve the right problems? The clearest way:

- **Consolidate** your data into a warehouse or FP&A platform that ingests everything, so that you have all the data available to your model.
- **Run an AI model** on this data warehouse so that it oversees the entire company.
- **Relegate internal tool AI use (as in Salesforce AI)** to answer specific operational questions where consolidation isn’t necessary.

Of course, there are also risks to letting an AI run across all your data, but we will discuss that in our Chapter 5 of this guide.

Consolidate your data into a single platform that ingests everything. AI only works when it can see across the full business.

Prioritize a small number of high-impact, repetitive problems where AI can deliver clear, measurable wins.

Start with fast, visible pilots to prove value, build trust, and set the stage for broader adoption.

Let’s get to work.

FP&A AI Sweetspots

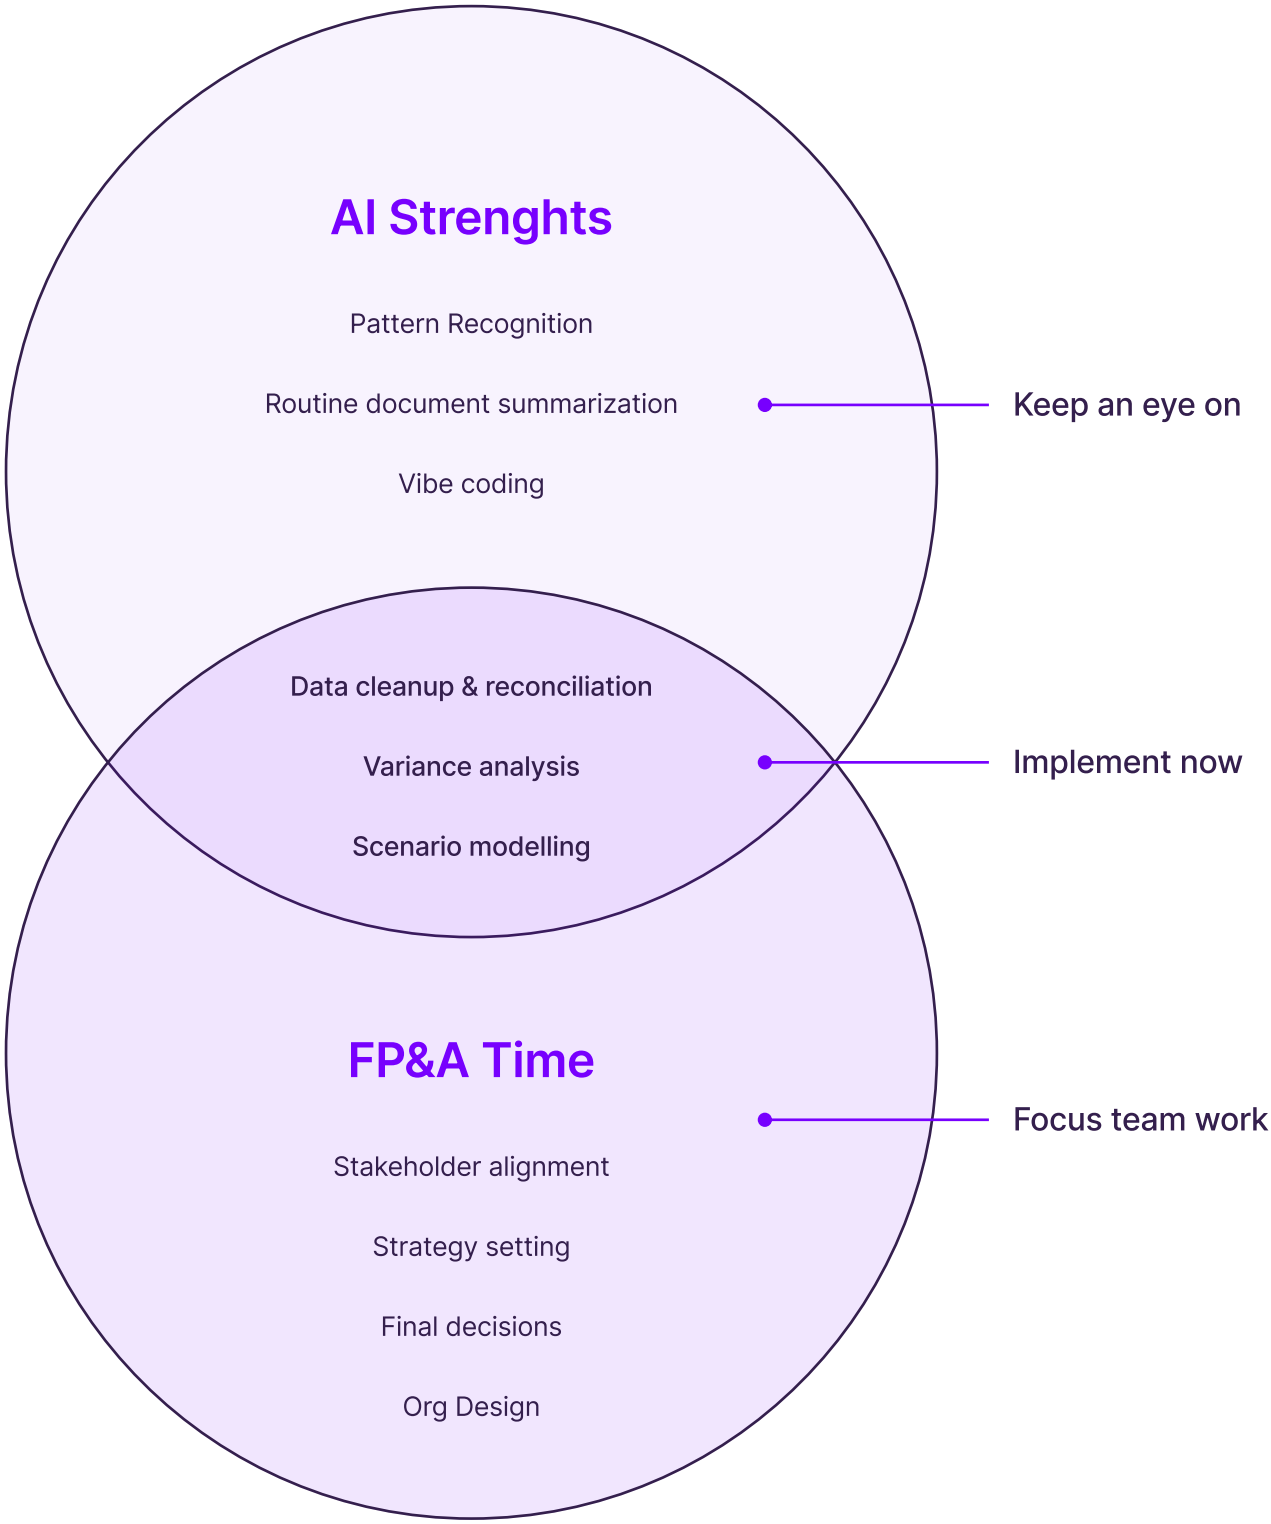

1 Consolidate Data and Questions

AI is powerful when it runs across a lot of data, not when it’s siloed in five different corners of the business. You wouldn’t want five different teams setting up custom data warehouses, right? So why would you want Salesforce, HRIS, Gong, and NetSuite all spitting out their own “AI answers” that you now have to reconcile with your central data warehouse? It’s hard enough to do that in reports. To get real value, you need to start by mapping out a few key things:

- **Explain what questions are answered where.** AI agents are now embedded in nearly every system from Notion to Salesforce to NetSuite. Your team needs a clear policy: What’s handled by these localized tools? What should be routed to your centralized AI implementation that sees everything? Which should have a dedicated report and which should be answered by chat?
- **Show what data is where.** At a detailed level. Pulling CRM data into your warehouse is great, but does that include customer success call transcripts? Sales emails? If not, where does that data live, and how can it be integrated? If you want to answer the hard questions, get everything you can in.
- **Know how the data moves.** If data starts in your CRM, when and how does it land in your warehouse? Is it a real-time sync, nightly batch job, or manual export? Knowing the pipeline flow is critical for data accuracy reports and transparency on when things should be correct.
- **Map into a clear process.** Make it clear to anyone how this all works through a simple data diagram that you present across to the teams. This gives a clear visual reference for how data is connected and where to find the answers to questions.

Centralizing your data into one granular, unified overview isn’t optional. It’s essential. That’s where AI actually becomes useful. Don’t hesitate to push everything possible (transcripts, meeting notes, etc.).

2 Prioritize Problems to Solve

AI is only as useful as the problems it’s asked to solve. So, start with just a few high-impact, solvable issues that show off AI’s strengths. How to do so:

- **Identify your issues.** Look for repetitive tasks that eat up time or require excessive attention to detail. Think data cleaning, creating different report versions from the same core set of numbers, or drafting initial responses to inbound requests.
- **Calculate the projected savings.** Talk to the people doing the work. Don’t assume that just because a process involved a copy-paste, automating it will save more than a few seconds. You need real numbers on time spent.
- **Analyze underlying tasks.** If the work is still manual, why? Is the data incomplete? Is it scattered across systems? And yes, AI should help you fix this by scanning across your data and suggesting the correct fixes.
- **Assess the ease of implementation.** Some tasks, like rolling up 10 subsidiaries in 15 countries, might sound exciting but are tough to automate. Go for the low-hanging fruit. Automate recurring reports. Set up a single successful chatbot or alert system.
- **Make the list.** Now that you have surveyed the scene, put the roadmap together.

The projects you choose will define how the rest of the company views AI. If your first few deliver only a couple hours of time-savings, skepticism will skyrocket. If they succeed, adoption accelerates.

Pick early wins that show the ‘wow’ factor for your implementations.

Here’s a framework to evaluate potential AI use cases:

Finance Problem	Wins	How	Time saved	Ideal examples
Monthly reporting package assembly	High	High	Medium	Clean GL + departmental input data
Headcount variance explanations	Medium	High	Easy	HRIS + budget vs. actuals mapping
Forecast version comparisons	Medium	Medium	Easy	Structured forecast versions
Ad hoc P&L analysis for execs	High	High	Medium	Granular P&L by entity/ cost center
Cash flow trend diagnostics	Medium	Medium	Medium	Historical cash flows + AP/AR aging
Financial policy Q&A (chatbot)	Low	Medium	Easy	Uploaded policy docs + chart logic

3 Centralize the team

You don’t need to hire a battalion of machine learning PhDs for internal AI use. They’re rare, expensive, and often not what’s actually needed. What you do need are sharp operators. People who can implement tools, integrate systems, map data flows, and train others. To do so:

- **Centralize AI learning.** Treat AI like any other enterprise tool. You don’t let 10 departments each figure out their own ERP strategy, so why would you let everyone run wild with AI? Build one core team that becomes your internal AI experts, trains others, and captures learnings.
- **Run it through Finance.** We’ve said it before: centralize tools and data under Finance. And, where the data lives, the AI should run. Finance is uniquely positioned to take charge. It already leads implementations, understands the systems, and knows how to train business users.
- **It’s a tool, not a build.** Don’t waste time trying to build proprietary models from scratch. The market has already spent hundreds of billions developing powerful, flexible AI tools. Your job is to plug them in where they add value.
- **Leverage executive and board support.** This isn’t hard to sell, especially when you pair your hands-on team with a board member or advisor who’s seen AI deployed successfully elsewhere. That outside validation adds credibility and speeds up buy-in.

You want your team to know how to use AI, where it fits, and how to deploy it, just like they would any reporting or planning tool

> Treat AI the same way you treat other tools: central ownership, business-led implementation, and ongoing training. That’s how you get real adoption.

4 Communicate Openly and Often

We’ve touched on this before, but it deserves its own spotlight because it can make or break your rollout. People are nervous about AI. It brings up concerns about job security, relevance, and change. And the best antidote is proactive, transparent communication. To do it right:

- **Set realistic expectations.** AI isn’t magic. Be upfront that there will be a learning curve, iterations, and some trial and error. Position it as a tool that improves and gets better over time.
- **Explain the “Why.”** Be clear about why the company is investing in AI. It is a tool, not a replacement. ERPs didn’t replace accountants. Emphasize how it will eliminate low-value work and create space for strategic thinking, analysis, and business partnership.
- **Be specific about impact.** Avoid vague promises. Lay out exactly which processes are changing and how roles might evolve. Highlight opportunities for upskilling and taking on higher-value tasks, not just process automation.
- **Show, don’t just tell.** When you land early wins (see next section), share them widely. Demonstrate how AI made someone’s job easier, saved hours, or helped close the books faster. Real stories will drive real buy-in.

Ignore the human element, and even the best AI tools will get quietly sidelined. Engage the team early and often, and you’ll turn skeptics into advocates.

> Over-communicate at the start. Regular updates, live demos, Q&A sessions, and executive sponsorship will build trust and momentum.

5 Bring it all together in a pilot or two

Like any significant initiative, AI adoption should begin with focused pilot projects. These pilots should directly tackle the high-priority problems that dovetail AI’s strengths with the organization’s needs. Some examples:

Pilot	Problem	AI Application	Success Looks Like
All systems source of truth	Data inconsistencies across systems are difficult to catch and correct. A customer disappears from CRM but still shows up in billing. A new deal closes but doesn’t show in HubSpot.	An AI agent runs hourly across CRM, billing, ERP, and CS tools to detect data mismatches.	The CRM logs a $20K ARR churn. The agent checks billing and ERP to confirm the drop. It identifies the churn as outside the normal range, flags a Slack channel with recommended actions, and confirms the issue once resolved.
Automated month-end variance analysis	Finance teams waste time gathering data and combing through dozens of views to identify budget variances.	An agent runs at month-end, analyzes key movements in the financials, and generates a summary with key drivers and supporting reports.	“MRR missed by 25% due to a 5% drop in close rate, driven by increased competitor mentions on calls. Usage-based revenue dropped 10% among top 3 customers due to seasonality. Additionally, customer calls declined 20% due to an unexpected employee absence.”
Executive daily briefing automation	Executives lack a consistent daily view of business performance, making it harder to stay aligned and proactive.	AI aggregates KPIs from sales, finance, operations, and customer service, flags significant changes, and compiles a daily executive briefing	By 8 AM each day, executives receive a concise summary highlighting 3–5 critical metrics or changes requiring attention, with links to relevant data.

Choose pilot projects that are visible, achievable in 1–2 months, and clearly tied to business outcomes. Every win sets the stage for broader adoption and deeper investment in AI.

In conclusion

To get real value from AI, you need structure and a strong roll-out. Those both start with defining how you’ll use AI in your organization and building the right foundation to support it.

Centralize your data. Centralize the team. Identify the right problems. And then pilot smart, fast solutions.

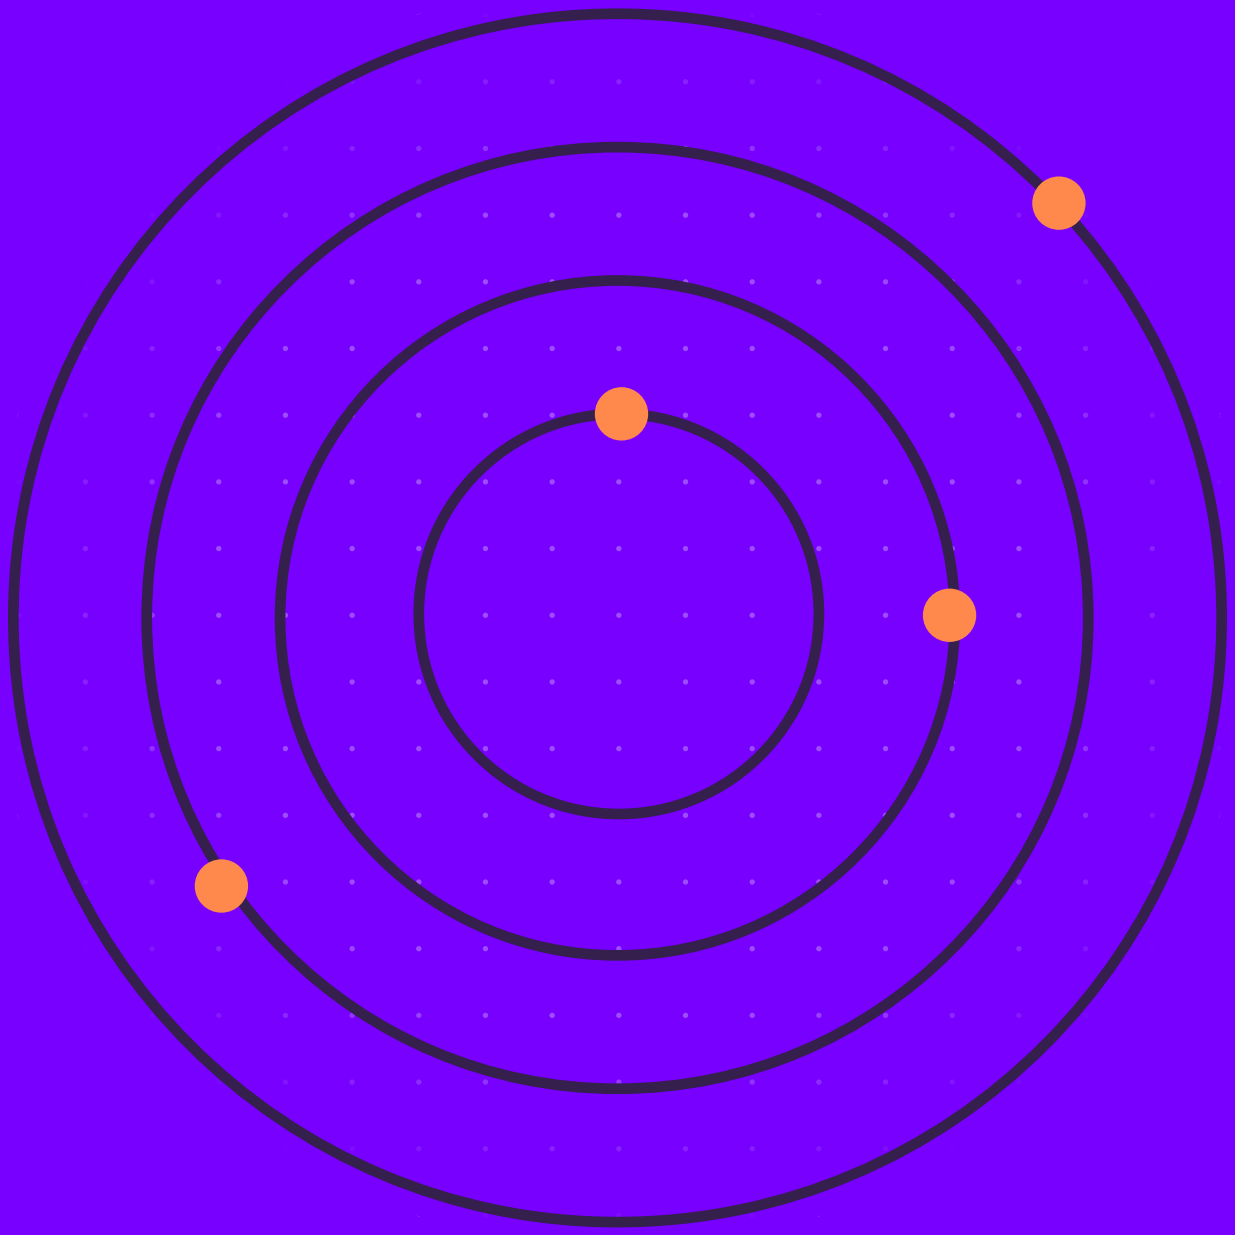

3. Why Finance needs to take charge

Finance should run with AI implementation

Finance teams should run all relevant data teams within an organization. We've talked about this at length, but simply put, Finance is best placed to run data teams, RevOps teams, and FP&A teams (of course) well. This becomes even more of a priority when you factor in the benefit that AI gains from centralized data.

Why? Because Finance has the right incentives and expertise to pull it off. Our job is to deliver actionable insights, drive accountability, and support decision-making across the business. And now that AI is becoming a meaningful layer on top of your data stack, the value of centralization has only grown.

Let's take a story from a couple of months ago when we share the example of someone who started leading finance at a tech company and found that employees were relentlessly confused: RevOps was buried in the CRM, Finance was stuck in the ERP, and the data team tried to reconcile everything in a data warehouse. The result? Everyone was confused, and any decision 'based' on data was spent resolving which information was right.

This will recur with any AI implementation if you haven't centralized your data team under Finance. With AI, it will just result in more of the same:

- **RevOps** siloed on 'responding to tickets' or pipeline projections without looking at how customers perform after they've joined
- **Data teams** siloed on rolling out infrastructure reports disconnected from business value.
- **Finance teams** generating countless new reports that don't integrate key non-financial metrics like usage, pipeline, or campaigns

But even worse, imagine letting the black box of AI work on all this data and then providing various “answers.” A decentralized approach will fail. It just adds three more sources for possible answers without solving the root problems.

Instead, you should centralize the effort into one one team, that owns the source of truth, that uses AI automations to look across all of your data, so you can make smarter decisions by acting on what AI does best: summarizing lots of data. Finance is this team for the same reason it should already run the data function. In terms of how to structure the team, I’d check this out (Build your FP&A team in the world of automated agents).

- All data converges to Finance, who is most incentivized to make sure that AI is using it and delivering the correct answers.
- Finance already has a strong cross-over between technical expertise and business strategy
- The team has the practice implementing complex toolsets and the incentive to automate the mundane part of their work

Let’s get to work.

1 Finance owns data

It's the core mission of the team: use data, in whatever form, to provide clear advice that leads to better decisions. In practice, this means combining operational data with financial data, the sweet spot of FP&A.

- **ERP, CRM, HRIS, and billing systems all converge in Finance.** No other team has this broad overview, nor the incentive to ensure it's always reconciled. (Example: reconciling headcount in HRIS with Sales CAC)
- **With Finance in control, there's one authoritative source.** The last thing you want is five different figures for ARR or conflicting views on pipelines based on unreconciled numbers. This will confuse everyone.
- **Ownership ensures consistent definitions** (revenue, customer, churn). And yes, this is more art than science, but if no one is overseeing the data that aligns from boardroom to front-line teams, you're effectively working from different books.

Finance often spends a significant amount of time just making sure all the data ties together. So lean into centralizing the team to make this work.

2 Finance is cross between the technical and business sides

This has always been Finance's job: converting numbers into insights. With AI, that role becomes even more critical by knowing how to check AI outputs (what's right or wrong), how to ask the right questions, and how to standardize those questions across the organization. This makes sense because the team is:

- **Already power users of data.** Finance has more power users of data than most teams. It's the people who write performance reviews in Excel and who see every table as a potential SUMIF.
- **Natural business partners.** Finance already supports other departments through budgets, tool evaluations, hiring plans, and more. The connections are already in place.
- **A viewer of the full picture.** The team understands how all parts of the company operate together across the customer journey in sales, marketing, product, support, and beyond.
- **Accountable to profitability, cash, and compliance.** These are the metrics that ultimately matter, and Finance is directly measured against them, making it even more crucial that these tools actually add value.

If you want a team to implement AI tools and use data well, that's literally the sweet spot of Finance: translating numbers into actionable advice people can actually use.

> AI will get things wrong, just like analysts. But Finance already has the muscle memory to quickly sanity-check numbers and know when something feels off.

3 Has a lot to gain from implementing internally before rolling out to others

We've hinted at this in previous projects, but it's worth stating outright: Finance has a lot to gain from implementing AI within its own team, so that they can be both the implementers and the users, creating a tight feedback loop.

Examples: Data cleanups, variance analysis, report building, scenarios, etc. This means:

- **Finance has motivation to roll out to save themselves time.** Finance is motivated to roll this out quickly so they can spend more time themselves on strategic insights instead of data reconciliation, which is the bane of the team's existence.
- **The team can validate the results.** If the team implements it *and has to use it* they gain tremendous insights into how much AI is worth it.
- **Able to share value across teams.** Quick wins in Finance (e.g., setting up a persistent data reconciliation system) can be taught and case-studied when implemented for other teams.
- **Offloading work efficiently.** Many Finance questions are about helping others understand how the numbers connect. If Finance implements it on their own numbers, they can invite others to ask questions which can be their first experience in AI to build trust.

By implementing AI within their own team, Finance can introduce the work to others and evaluate if it ***actually*** delivers ROI. Win-Win.

> Yes, Engineering teams are using AI to speed up coding, but that's a highly technical use case which isn't as applicable to other teams that aren't coding.

4 Has practice implementing tools for teams

In today's modern tech stack, the key is integration to make sure the data actually connects across systems. Finance is already in the thick of this, often acting as the glue between tools and teams. This makes Finance a natural fit for leading AI implementation across departments because of its:

- **Experience in ERP rollouts.** Finance often leads (or shoulders) ERP implementations. They're used to complicated rollouts and know how to push through the pain points.
- **Role in tech stack integration.** Finance lives and dies by its tech stack. When integrations break, reports don't run. That makes them vigilant people who care about getting the plumbing right.
- **Service of other departments.** Finance is used to supporting other teams. That same mindset applies when implementing AI that answers knowledge-base questions or automates repetitive queries.
- **Skepticism of tools.** Let's be honest, Finance leaders have seen every flavor of over-promised software, which makes them ideal to cut through the hype and buy tools that actually deliver.

Implementing AI across a centralized tech stack isn't easy. But if you want your company set up for success, Finance is the team that knows how to do the heavy lifting and see it through.

In conclusion

Implementing AI across a centralized tech stack isn't easy.

That's why you want the rollout led by the team best incentivized for success. Finance brings the right mix of experience, cross-functional perspective, and incentive alignment.

They will be focused on delivering tools the business will actually use to make smarter, faster decisions.

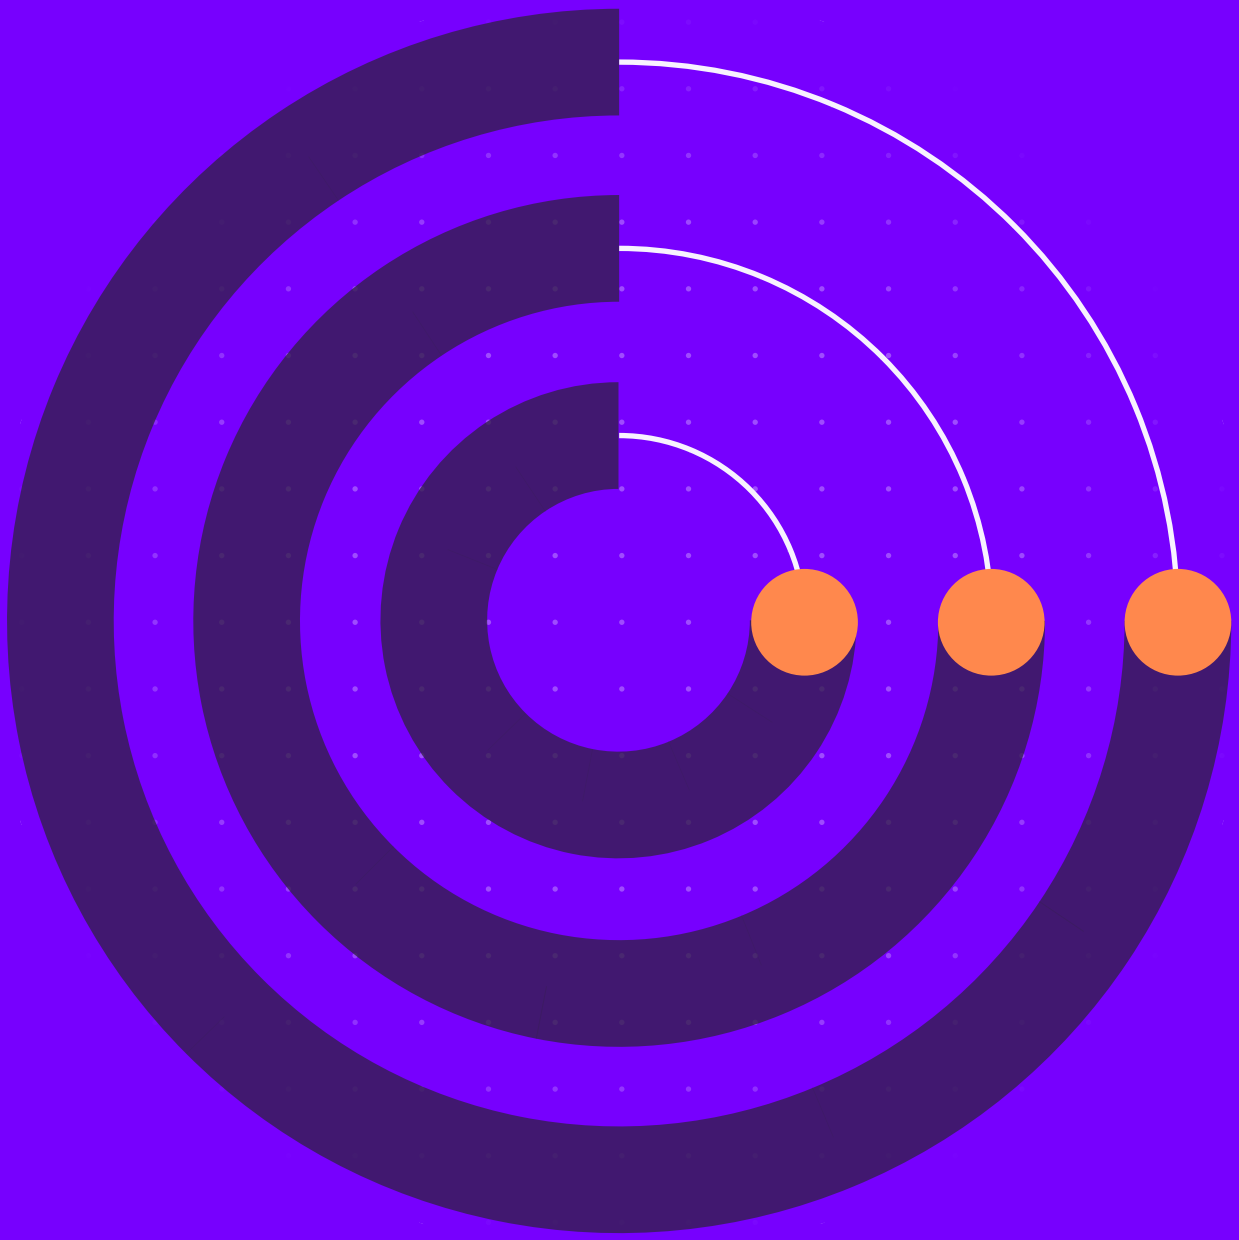

4. Monitoring Performance & ROI

The rollout is 10% of the battle

Implementing an AI tool is just the first step. The novel part of AI is how quickly it can learn and adapt, especially as your team figures out how to interact with it. But that learning doesn't happen automatically. It needs to be guided.

It's been a while since our last story, so let's start with one.

We had spent months and months implementing a new ERP system. The goal? Get our inventory stock levels to properly integrate with our financials. Pretty standard stuff. We brought in consultants, ran endless training sessions, handed out shiny new tablets for the warehouse teams. Nothing wildly complex, but we had a lot of locations with multiple warehouses. Accountants had spent six months mapping data, reconciling accounts and loading everything into the new system.

Then came go-live day. We flipped the switch... and things went quiet.

Given the nature of our business, a lull was expected. We weren't moving inventory at a crazy pace, so we patted ourselves on the back and shifted focus to other projects, like budgeting. We figured the system was up and running and the numbers looked solid due to the automated reconciliations. So, we left it alone.

Fast forward a few months to our next major distribution cycle and red flags started flying. It turned out, staff on the ground had just trusted the system to work without review. They had used an "auto-reconcile" function to bypass any actual stock counts and moved on. It worked... until we got to audit prep and found our inventory was off by nearly 20% with ins and outs that looked insane. And across product lines? It was chaos.

The core issue? We had exhausted ourselves getting to go-live and then let the follow-through slip. The inventory teams never saw a balance sheet so they had no trigger to fix things.

So yes, everything we've covered so far in terms of consolidating data, prioritizing problems, centralizing teams set the stage. But real, lasting payoff comes after go-live, when you are making smarter, faster decisions.

Practice / Focus Area	What to Measure (KPI)	Quick Tip
1. **Define Success Metrics Up-Front**	Days-to-Close, Retention Rate, Employee-to-ARR, Forecast Accuracy	Let an AI agent auto-pull and format the KPI pack.
2. **Analyze AI Use**	Active users, time / user, Feature utilization, AI vs. legacy usage	Cull "dead" reports to spotlight real adoption.
3. **Review & Update**	Active users, time / user, Feature utilization, AI vs. legacy usage	Define "kill criteria" early to avoid zombie projects.
4. **Data Quality Scorecards**	Open data-quality issues	Define "kill criteria" early to avoid zombie projects.

First, clearly define the success metrics up-front that you are aiming to achieve.

Make sure you are identifying high-impact business problems.

Launch focused pilot projects on low-hanging fruit to deliver quick, tangible wins and build crucial buy-in for your AI strategy.

Let's get to work.

1 Define Success Metrics Up-Front (and Stick to Them)

Ideally, this happens before you even think about implementation. But as the AI starts running, you must quantify the impact you expect and track it. It sounds obvious, but this step often gets overlooked. Done right, it builds credibility, reveals what's working (or not), and sharpens future rollouts.

- **Translate outcomes into hard KPIs.** Focus on business results. Not just activity metrics or technical stats. Some examples:
  - → **Days to Close Books:** Direct measure of Finance efficiency.
  - → **Retention rates:** You should take a cohort of customers that you are piloting AI with, in say CS, to measure the retention.
  - → **Employee → ARR:** You should see an uptick in scale as employees are able to deal with more.
  - → **Forecast Accuracy %** (Sales, Demand, Financial): Tells you if predictive models are improving.
  - → **Cost per Invoice Processed:** Shows real operational savings.
- **Gather baselines.** Know exactly where things stood before the AI rollout. Use hard numbers, not estimates or gut feel.
- **Make AI KPIs part of MBRs/QBRs.** These aren't just side metrics. Bake them into your regular business reviews. Have an agent itself pull, format, and present these metrics and make sure they are correct.
- **Track time saved, carefully.** This one's trickier but still essential. Consider short surveys, time studies, or calendar audits to measure how your team's workload shifts. Are they spending less time cleaning up spreadsheets and more time on strategic analysis?

By setting these standards clearly, based on tangible outcomes, you'll ensure everyone is laser-focused on delivering successful results, not just "implementing AI."

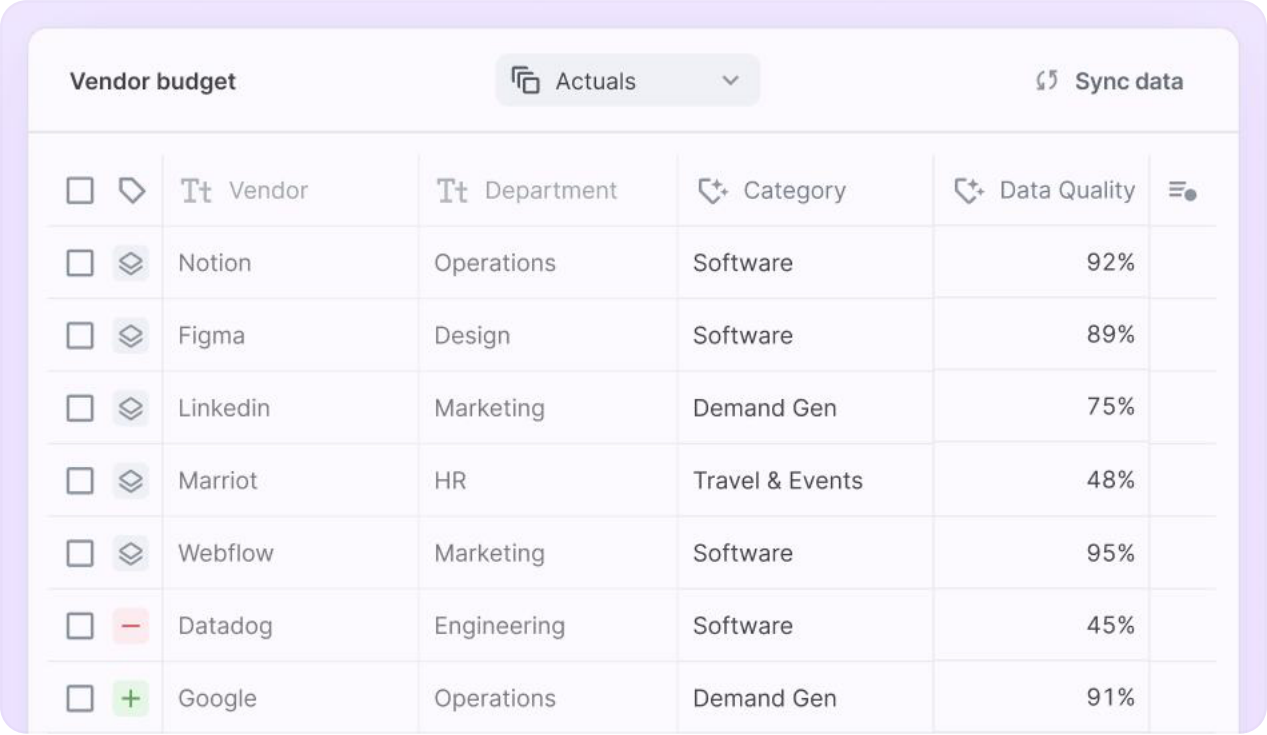

Vendor budget | Actuals | Sync data

Vendor	Department	Category	Data Quality
Notion	Operations	Software	92%
Figma	Design	Software	89%
Linkedin	Marketing	Demand Gen	75%
Marriot	HR	Travel & Events	48%
Webflow	Marketing	Software	95%
Datadog	Engineering	Software	45%
Google	Operations	Demand Gen	91%

Create data quality scorecards that assign a grade (A, B, C) or a percentage score to the key datasets feeding your AI. Review and update these monthly. If garbage goes in, garbage comes out, no matter how smart the AI.

2 Effective Training

AI skepticism is real and understandable. Training should be about removing the mystery, building confidence, and showing teams how to work with AI. Poor knowledge sharing is a guaranteed path to low adoption and wasted investment.

- **It's a mindset shift, not just a skillset.** Frame AI as a partner. Focus on how it takes work nobody likes off their plates so they can spend more time on strategy, analysis, and creative problem-solving.
- **Make knowledge sharing role-specific, not one-size-fits-all.** Your teams don't need a tutorial on prompt engineering. They need to understand strategic implications and how to ask the right business questions of AI. Analysts need to grasp inputs, assumptions, and outputs. Front-line users need to see exactly how AI fits into their daily workflow.
- **Build champions and make learning ongoing.** Identify early adopters and empower them to be internal super-users. Peer learning spreads faster than formal sessions.
- **Use AI to answer 'how.'** The biggest benefit of a chat-based interface is that it's dead simple. People should ask questions at their own pace and dig down where they need to.

For more on this, check out how to build a finance team in the world of AI agents.

> Encourage the use of AI in management reviews, but no need for formal training programs. You need to change the mindset, which is tough to do in formal training systems.

3 Analyze AI Use: Are They Using It, and Is It Working?

While outcomes are king, usage is the leading indicator. Measure its use. Of course, high usage doesn't automatically equal high value, but when you see usage climbing alongside improved business outcomes, it's a strong story. What to track:

- **Accuracy.** Track how often the AI implementation is right. You can do this by sampling questions against reports that you already trust or asking everyone to do this for the first month and giving a like/dislike response to the answers they get.
- **User adoption count.** Track how many intended users have actively started using the AI system (e.g., "15 out of 20 accountants ran at least one AI-driven forecasting report in the first month").
- **Frequency of use.** Measure the number of AI sessions, queries processed, models run, or reports generated per user/team per week/month.
- **Feature utilization.** If your AI tool has multiple features (e.g., forecasting, anomaly detection, natural language querying, report generation), track which ones are being used most and by whom.
- **AI Usage vs. legacy.** This is crucial. If possible, compare the usage of the AI tool against the old method it was meant to replace. A proxy might be to monitor the "number of manually created forecast spreadsheets submitted" going down as AI forecasting usage goes up.

If you track this at an individual employee level, you can start to correlate AI usage with performance on specific tasks. This allows you to not only prove the AI is working but also to identify power users who can champion the tool and help others get up to speed.

> Get rid of dead reports. Things that were created but aren't being used. It's probably for a reason and creates a tremendous amount of headache.

4 Review and update

Once your AI setup is live, it needs ongoing attention. Regular reviews ensure it becomes even more valuable as time goes on.

- **Monthly operational huddles.** Focus on usage and functionality. Are people using the tool? What's broken or needs tuning? Review adoption, usage trends, data quality, and early KPI wins. Keep it tactical, fast, and focused on immediate improvements.
- **Quarterly strategic reviews.** Step back and assess business impact. Review outcome-based KPIs against your original goals. Is AI delivering material ROI? What new opportunities has it unlocked? Adjust your roadmap and resource plan accordingly.
- **Know when to walk away.** Set clear kill criteria upfront. Low adoption, no measurable impact, or rising costs with no return? Time to sunset. Don't let zombie projects drain momentum.

This is nothing different from any other project. Assess how the rollout is doing, plan the next steps.

> Like a broken record... but let the AI draft its own quarterly performance report. Use it to guide the discussion and demonstrate its capabilities.

5 Keep up to date

AI will continue to evolve quickly. Stay current to maintain your competitive edge and take advantage of the significant roll-outs that are happening, at this point, about once a quarter.

- **Tap into peer networks and communities.** Encourage your team to join industry groups, webinars, and AI communities. Hearing what's working (or failing) for others accelerates your own progress.
- **Make learning part of the culture.** Teams using it need to treat learning and adaptation as part of the job. Bake continuous improvement into the way you work.
- **Assign ownership to track what's next.** Make sure your centralized team is periodically looking for new use cases once you have your data consolidated.
- **Get the most from vendors.** Your vendors are building these tools. Push for roadmaps and new feature access. Hold them accountable to help you stay ahead.
- **Balance innovation with ROI.** Stay current but don't chase every shiny AI upgrade. Vet new initiatives like any investment: strategic fit, value delivered, and cost to implement. Innovation is only useful if it moves the needle.

This technology is moving faster than previous 'tool' upgrades which usually came yearly. Now, you should expect improvements at least each quarter.

> Your problems should send you searching for technology that will fix it, not for technologies searching for a solution.

In conclusion

Getting AI live is just the start.
The real value comes from disciplined follow-through. You need to define success with hard KPIs, track usage closely, and regularly review what's working and what's not. Train your team to partner with AI, not just use it, and stay current as the technology evolves. The edge goes to companies that treat AI as a capability to grow.

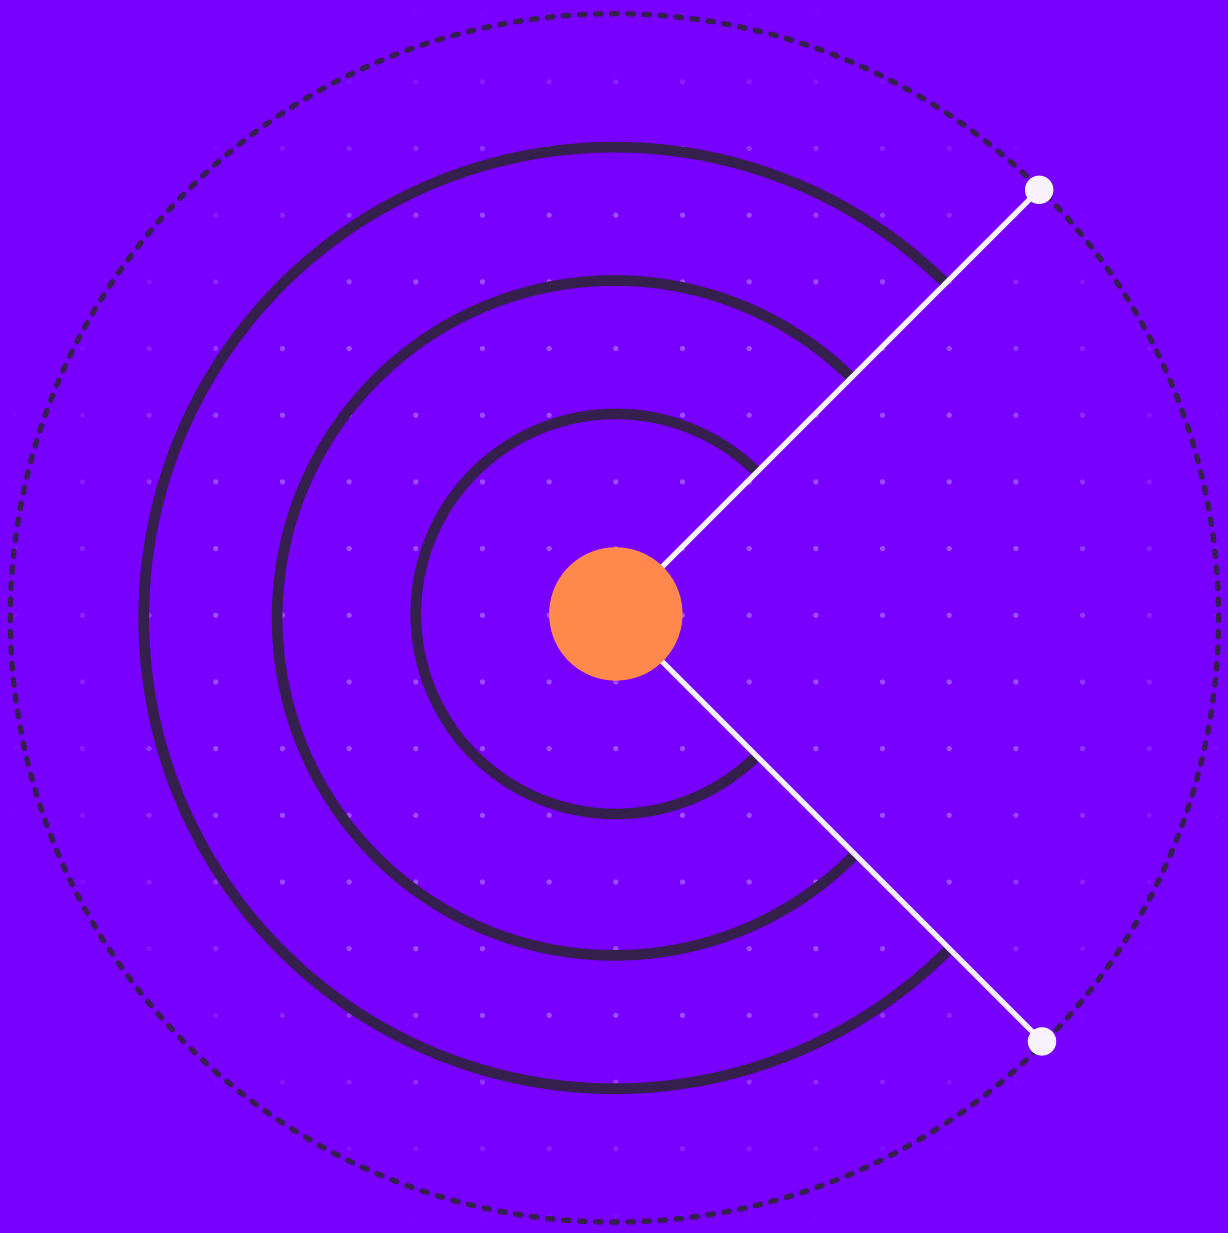

5. AI regulations and risk. Fun times.

Making sure it all works

While it's much more fun to talk about moving fast, boosting productivity, and solving all your problems with AI, it's worth slowing down to talk about the risks. Whether you're implementing a tool or building something yourself, the dangers are real. Of course, you don't hear about them as much as there's a lot more money in promises.

These risks fall into three buckets.

1. Wrong Answers	"Sounds good" isn't good enough for finance. The team's value comes by delivering objective guidance based on data. If the data or analysis is wrong, the team actively hurts the company.
2. Internal Data Leakage	You don't want someone asking a chatbot, "What is my manager's salary" or "Can you give me all the key data on our biggest clients so that I can take it to our competitors?" and receiving answers. At the same time, you do want your leaders to see across all data and provide insights. Fine-grained permissions are key to making this work and a lot of AI hasn't reached that maturity at this point.
3. Regulatory and data exposure	There are compliance concerns, too. In Europe, GDPR still reigns supreme. In the U.S., regulation is catching up fast. It is easiest to comply with the stricter set of laws (EU) to ensure you can deliver to any clients working in both markets.

In the worst-case scenario, all three risks hit at once: AI is unleashed across all your data, people start getting wrong answers that confuse decision-making, sensitive internal information leaks, and suddenly your customers are seeing one another's data. Now you're dealing with internal chaos and external regulatory fallout. Prompt injections and clever work arounds have already been successful getting supposedly sensitive information shown in general models.

The best way to avoid this is to focus on prioritizing projects where you can easily check the results and control your permissions (hello Finance). That means initial rollouts like: cleaning data and creating forecasts you can validate.

AI isn't "smart" in the way we're used to. A human who has a PhD can probably sum up numbers correctly. Whereas AI can deliver a PhD-level market analysis one minute, then miscount how many "r"s are in "strawberry" the next. It's deeply unintuitive. That's why it's crucial to double-check any AI rollout.

There are three major types of risk:

- Wrong answers that cause confusion.
- Internal data leakage within the company
- Regulatory exposure.

Regulation is coming. The easiest way to prepare is to hold yourself to GDPR standards across both the US and Europe to ensure there are no issues. Aka: transparency, opt-outs, private information stays hidden.

When you are considering building vs. buying, you need to critically assess each one of these risks.

Let's get to work.

1 Sounds nice but...

AI is built to sound convincing but that is not the same as being right. It is notoriously bad at trying to cover up its mistakes by lying without regret. On the plus side, a big part of finance is built to cut through people trying to cover up bad numbers. On the downside, other teams aren't so well trained. Here are the ways it can go astray:

- **Wrong answers.** The whole point of using AI is to get better data into more hands, faster. But when the answers are wrong, the decisions that follow are wrong too and the impact is multiplied.
- **Wrong questions.** And, no this isn't solved with prompt training. Asking a good question is tricky even in the best of times. Many non-finance users don't understand key metrics, time frames, or inclusions/exclusions in the data. If the question is off, the model interprets it incorrectly and now you've got a mismatch before you even get an answer.
- **Explainability.** In a report, you can trace the data. You can see the rows, the formulas, the assumptions. This is hard to do with AI. Understanding the actual calculation behind a chatbot's answer is often nearly impossible, making it tough to verify anything.
- **Confirmation bias.** AI models are people-pleasers. Ask, "Is my performance strong?" and they'll often try to say yes. But the whole point of having numbers is to challenge assumptions, not reinforce them.
- **So easy...** It's so easy just to ask a question, take an answer, assume it's generally right, and move on. This behavior is reinforced by how people interact with general-purpose AI models every day. In an organization, this becomes dangerous fast when the only rationale for decisions is, "That's what the AI told me."

Getting data to be right even 90% of the time (much less 100%) using AI models is challenging even for the biggest companies in the world burning billions of dollars in pursuit of that truth. Be wary of these challenges as you test rollouts.

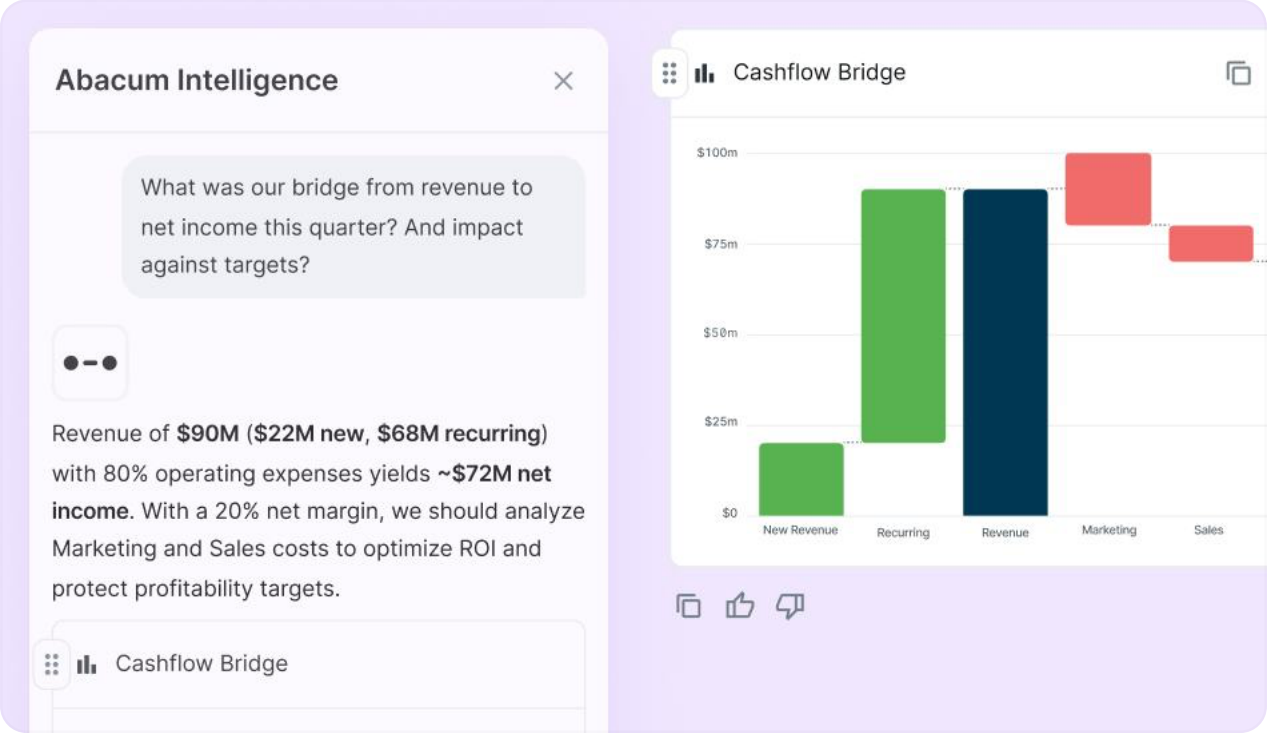

Prompt training is not the solution. You need to continuously check and monitor the performance by looking at actual chat requests and answers and then comparing it with trusted reports. You also need to establish that each person has ownership of their decisions no matter what the AI says.

2 Data leakage

Just like with any sensitive data system, you need firm access control and data security. That means tightly limiting the subset of data AI can access: per person, per use. It's not enough to trust that users will ask the "right" questions or that the system will always interpret them correctly. Additionally, prompt injections and clever workarounds already exist in general-purpose models. Left unchecked, these loopholes can lead to intended or unintended breaches.

Here's what you risk sharing:

- **Private company information.** If you open up a chat interface to your team, people will inevitably ask curious, probing questions. Without proper controls, sensitive data like salary details, board materials, contracts, or your cap table could become widely accessible.
- **Private client data.** The same applies to customer records: documents, contracts, personal data. If a model can access it, someone might find a way to extract it and use it against the company.
- **External access risk.** More and more companies are embedding AI into client-facing dashboards. Clients expect visibility, but the last thing you want is for one client to start pulling data on another. That's how you lose trust, fast.

The bottom line: AI isn't a person. It doesn't understand boundaries, and it can be manipulated. That's why whatever you build must have strict, role-based permissions and deep access control built in from the start.

Pressure test this in any way you deploy AI.

3 Regulation US and Europe

There is a lot coming down the pipeline, and this is **not** legal advice. But generally, best practice is to follow the same policies in the US as you would in Europe to future-proof yourself from regulatory action. Generally that is:

- **Know where high risk lurks (Especially in HR).** AI used in hiring, promotions, or performance evaluations falls under "high-risk" use cases. These require enhanced transparency and oversight, with compliance deadlines looming in 2026 in Europe.
- **Don't assume summarization equals anonymization.** Just because an AI tool summarizes data doesn't make it GDPR-compliant. Individuals must be told how their data is being used, even in summary form, and there must be meaningful human oversight when automated decisions impact people.
- **Respect the right to opt out.** The EU AI Act grants individuals the right to opt out of their data being used to train AI models. FP&A teams working with SaaS vendors need to ensure these rights are both contractually protected and technically enforced.

- **Be ready to disclose breaches, fast.** Significant AI-related incidents must be reported within 15 days under the AI Act and within 2 days if individual rights or safety are at serious risk. Finance teams should have protocols in place for rapid incident response in collaboration with legal and IT.
- **Data deletion isn't just hitting "delete".** GDPR and the AI Act empower individuals to request their data be erased, including from training datasets. For many models, this means retraining or re-engineering the system, something finance teams need to assess with vendors up front. This was the issue with ChatGPT being paused in Italy.
- **Watch where the data lives.** Data residency laws, particularly in the EU, restrict where personal data can be stored and processed. AI tools that pull in data from multiple regions must be compliant with these geographic constraints, especially when dealing with sensitive employee or customer information.

Again, these are the best practices. But, this isn't legal advice. It's just a quick checklist on the different types of expectations you should have as you run AI over your databases.

This usually becomes a very high priority as you scale or go through a funding round, which can cause a tremendous amount of re-done work, so it should be considered early on.

4 Build vs. Buy

Getting access control, data validation, compliance, and everything else is expensive and time-consuming. But it's also mission-critical. So whether you're building internally or buying from a vendor, you need to:

- **If You Use a Vendor...**
  - **Kick the Tires.** Don't take anything at face value. Dig into the details: data segregation, security incident response, uptime SLAs, permissioning logic. Make sure it's real and proven.
  - **Ask their clients.** Have they had any specific issues in this regard or how do they handle it internally?
  - **Contract Shared Risk.** A solid vendor contract will include liability provisions and compliance obligations. If something goes wrong, you're not carrying the entire burden alone.
- **If You're Building In-House...**
  - **Budget Realistically.** Building isn't just about development. You also need compliance, access management, testing, auditability, and ongoing maintenance. Don't underestimate the full scope.
  - **Plan for Scale.** Today, it might be fine if two finance analysts can see everything. But what happens when your org scales 5x? You need a long-term architecture, not a quick fix.

In short: yes, there can be strategic advantages to building in-house. But there are also real downsides. Whichever path you choose, treat the decision like any other major system investment and evaluate it thoroughly.

Make sure your vendor has staying power. Look for strong financial backing, long-term viability, and references from similar-sized customers. You don't want to be left hanging in year two.

In conclusion

Using AI well can have efficiency gains of up to 50-100%.

But, using it badly can have a tremendous negative impact on the company. So, spend the time and money to do it well.

In conclusion

In our second guide, we get into the specifics of actually using AI. We began with 1. What to expect from AI today to give an overview of what is possible. Then we got into actually 2. How to roll out "AI" properly before discussing 3. Monitoring Performance and ROI. We then ended with discussions on 4. Why Finance needs to take charge and 5. AI regulations and risk. Fun times.

We hope you enjoyed. Thank you for your time.

Abacum

Discover Abacum

Abacum's mission is to empower finance teams to work smarter, make better decisions, and accelerate their company's success. Our FP&A platform simplifies business planning to drive performance. In weeks (not months) you will have:

Single source of truth. Gain a complete, accurate view of your business by connecting all your systems into one platform.

Automated reporting. Consolidate your reporting and identify trends and variances with dynamic dashboards, enabling strategic actions.

Insights on every decision. Influence business outcomes with data-driven conversations and actionable insights.

Plans without limitations. Adapt to any changes in business plans, and easily scale your processes.

Abacum enables finance leaders to focus on delivering results, not reports. To see more, check out our detailed workflows for...

Budgeting & Forecasting | Revenue Planning | Headcount Planning | Investor Reporting | Financial Statements

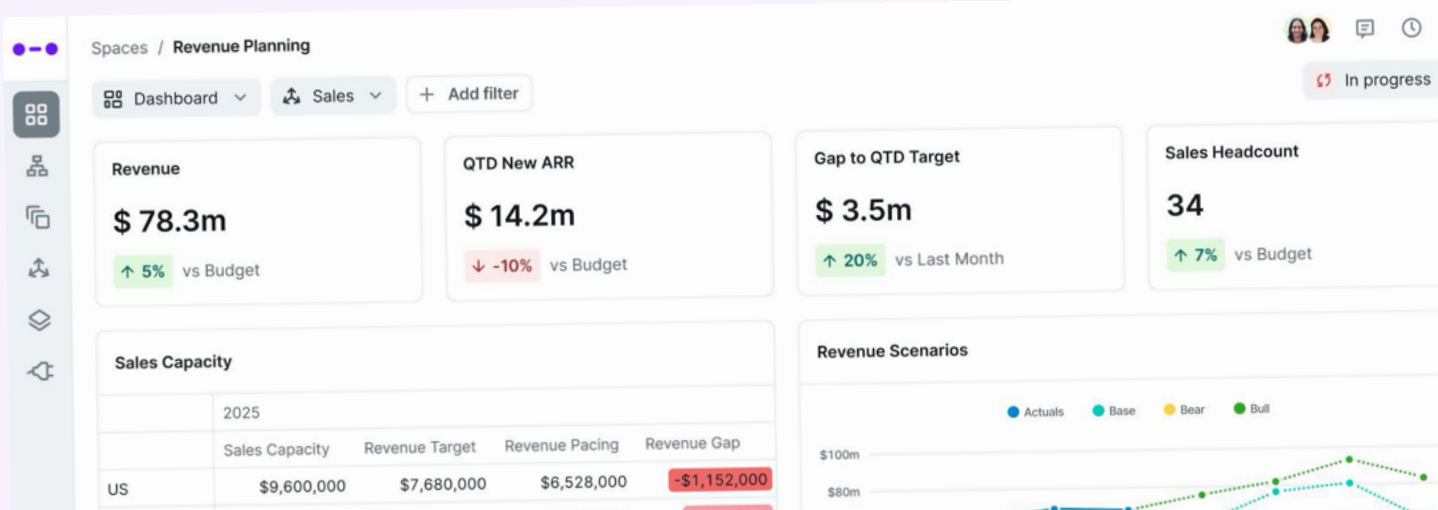

Abacum

Business Planning to
drive performance

Book a demo