How to do AI Cash Flow Forecasting

the ultimate guide for treasury & finance teams

APRIL 2026

VERSION 1.0

Nine Ants®
by Automation Boutique

Table of contents

PART I: THE FOUNDATION

Introduction

Every treasury team wants AI cash flow forecasting. Almost nobody is ready for it.

Imagine a mid-size manufacturer. Twelve subsidiaries across eight countries. Eight currencies. The CFO walks into the quarterly business review and asks: "Where will our cash be in thirteen weeks?"

What happens next is usually a spreadsheet nightmare. Treasury sends emails to every subsidiary requesting their forecast. Half respond on time. The data comes in different formats. Someone forgot to include a tax payment. Another subsidiary submitted numbers in the wrong currency. The team spends two weeks consolidating, cleaning, and reconciling. Only to produce a forecast that is already outdated by the time it reaches the CFO's desk.

This scenario plays out in thousands of companies every month. It has played out for decades. But something has changed. AI, and more specifically agentic AI, has matured to the point where it can genuinely transform how treasury teams forecast cash flows. As a practical tool that delivers measurable results.

At Automation Boutique, we have been building AI solutions for treasury and finance since the technology became mature enough to deliver real value. We have seen what works, what does not, and where most companies go wrong. The pattern is almost always the same: they jump straight to the AI part and skip everything that makes AI actually work.

AI cash flow forecasting is not magic. It is a system. A system that starts long before any model runs its first prediction. And if you get the foundation wrong, no amount of AI will save you.

This is the playbook. Step by step. From raw bank data to autonomous treasury operations. Whether you build it yourself, use spreadsheets, or use a platform like NineAnts, the principles are the same.

“

AI cash flow forecasting is not magic. It is a system. A system that starts long before any model runs its first prediction. And if you get the foundation wrong, no amount of AI will save you.

JAN-WILLEM ATTEVELT | CO-FOUNDER, AUTOMATION BOUTIQUE

THE BUSINESS CASE

Why Cash Flow Forecasting Matters

What is bad forecasting costing you right now? Calculate it honestly, and the business case for better forecasting writes itself. We help treasury teams put that number on the table, so they can measure results.

JAN-WILLEM ATTEVELT | CO-FOUNDER, AUTOMATION BOUTIQUE

Cash flow forecasting is consistently ranked as the number one challenge in corporate treasury. Not payments. Not bank relationship management. Not FX hedging. Forecasting.

Why? Because it is the foundation for almost every other treasury decision. And most companies are still doing it poorly.

What happens when you get it wrong

A company with inaccurate cash flow forecasting is essentially flying blind. The consequences are real and measurable.

Idle cash earns nothing. If you do not know that you will have 10 Million of surplus cash for the next sixty days, that money sits in a current account earning zero. At a conservative money market rate of 3%, that is 50,000 in missed yield in just sixty days. Multiply that across a year and the numbers become significant.

Emergency funding is expensive. When a cash shortfall catches you by surprise, you do not negotiate from strength. You draw on credit facilities at unfavourable terms, or worse, you scramble for overnight funding. The spread between planned and emergency funding can easily be 200 basis points. On a 20 Million shortfall lasting sixty days, that is 65,000 in unnecessary interest. For a single event. And that is only the external side of the story. With a reliable forecast, you can often smooth cash flows internally before going to the market at all: postponing non-critical payments, swapping liquidity across currencies, or using cash intercompany to cover a short-term gap. That kind of proactive management is only possible if you see the shortfall coming.

Unhedged FX exposure hits the P&L. If your forecast does not accurately capture the timing and size of cross-currency cash flows, you cannot hedge effectively. A 5% adverse move on EUR 10M of unhedged USD exposure is EUR 500,000 directly off your bottom line.

And perhaps most damaging: poor forecasting means poor visibility into business health. Cash is the ultimate truth-teller. When customer receipts start declining or supplier payments start stretching, these are early signals. Without a reliable forecast that tracks these patterns and compare them with previous years/seasons, management discovers problems too late.

What is bad forecasting costing you?

MISSED YIELD	EMERGENCY FUNDING	UNHEDGED FX
Idle balance **10M**	Unplanned shortfall **20M**	FX exposure **10M**
Money market rate **3%**	Spread over planned **2%**	Adverse FX move **5%**
Missed yield **50k** (per 60 days)	Cost per event **65k** (per 60 days)	P&L hit **500k**

Illustrative inputs for a mid-sized group. Plug in your own.

What happens when you get it right

The flip side is equally compelling. Companies with accurate, timely cash flow forecasts consistently achieve better investment yields, lower funding costs, tighter FX risk management, and earlier warning of business issues. Treasury stops being a reporting function and becomes a strategic partner.

The ROI calculation is straightforward. Take your average idle cash balance, multiply by the yield you could be earning, and add the cost of emergency funding events and FX losses over the past year. For most mid-to-large corporates, the annual value of improving forecast accuracy by even 20-30% runs into hundreds of thousands, sometimes millions, of euros or dollars.

If you are trying to convince your CFO to invest in AI forecasting, start there. The numbers speak for themselves.

Look Back Before You Go Forward

Before you forecast a single Dollar or Euro, you need to understand your past.

Historical transactions and bank account balances are the raw material of any cash flow forecast. Without them, you are guessing. With them, you are learning. This is the unsexy part that most companies want to skip. Do not skip it.

Get your bank data. All of it.

You should aim for at least 24 months of historical transaction data. Why 24 months? Because you need to see full annual cycles. Seasonality, the recurring patterns in your cash flows that follow a calendar rhythm, only reveals itself when you have enough history to compare year over year. Tax payments hit in specific months. Customer receipts might spike in Q4. Supplier payments might cluster around quarter-ends. You cannot detect any of this from three months of data.

If you can get 36 months or more, even better. More history makes it easier to distinguish true patterns from one-off anomalies. That unusually large payment in March 2025: was it a seasonal pattern or a one-time event? With three years of history, the answer becomes clear.

How to automate bank connectivity

Getting bank data into your system automatically is important. There are two main approaches that most corporates use:

Host-to-Host (H2H) via SFTP: the workhorse of corporate bank connectivity. Your bank drops statement files, typically in CAMT.053 or MT940 format, onto a secure SFTP server. Your system picks them up automatically. It is batch-based, not real-time, but it is reliable and well-established.

API connectivity: the modern approach. Banks increasingly offer APIs that let you pull transaction data in real-time or near-real-time. In Europe, PSD2 and Open Banking regulations have accelerated this. Many banks now offer standardised API access that did not exist five years ago. The data is fresher, the integration is cleaner, but API coverage is not yet universal across all banks and geographies.

For companies operating globally, SWIFT messaging adds another connectivity option, particularly useful when dealing with a large number of banking partners across different regions.

The reality for most multinationals is that you end up using a combination of these methods. Your German subsidiary can get data from Deutsche Bank with H2H file drops. Your US entity can get data from its JPMorgan bank account via API. That is not a technology problem. It is an integration problem. And it must be solved before any forecasting can happen.

If you are still downloading CSV files manually from your bank portals every month, stop. Invest in automating your bank data pipeline. Whether you build the integration yourself, use middleware, or use a platform like NineAnts, the goal is the same: connect via API or SFTP and normalise the data into a single format regardless of source. Automation here pays for itself almost immediately. CSV uploads can serve as a fallback, but automated feeds should always be the preferred path.

Categorise your cash flows

Raw transactions are noise. Categorised transactions are signal.

Define the cash flow categories that matter for your business.

The usual suspects (categories often used)

Customer Receipts	***Bank Fees***
Supplier Payments	***Interest***
Salaries	***Rent***
Taxes	***Utilities***
Capex	***Loan Drawdowns and Repayments***
Intercompany flows	***Dividends***

But do not just copy this template. Make them specific to how your business actually operates. A SaaS company and a manufacturing company have very different cash flow profiles.

Now comes the hard part: labelling every historical transaction with the right category.

There are a few ways to approach this. If you want to use what you already have, Excel's Power Query engine can handle rule-based categorisation, and it is completely free. You define rules: if a transaction has these characteristics, categorise it as X. Power Query applies them across your data. It works, but it requires a lot of rules, a lot of maintenance, and a lot of backtesting to get right.

A more scalable approach is a hybrid model: rule-based categorisation for the obvious patterns, with AI classification picking up everything the rules miss. This is the approach we took with NineAnts, but the same principle applies whether you build it in-house or use a third-party tool. Rules catch the straightforward stuff. The payroll that comes from the same counterparty every month. The rent that is always the same amount. AI handles the ambiguity: the supplier payment that could be Capex or could be regular procurement, the bank charge that is structured differently across banking partners. The result is fully categorised transaction history without the manual grind.

Watch out for common categorisation pitfalls. Intercompany transactions are a frequent source of miscategorisation. A payment from your Dutch entity to your German entity might look like a regular supplier payment if you are not careful. Refunds and reversals can throw off your category totals if they are not matched to the original transaction. And multi-currency transactions require extra attention. The same supplier might appear with different names depending on which bank is processing the payment.

At this point, you have not used any AI for forecasting. And that is the point. You now have a clean, categorised view of your cash flow history and historical balances. This is the foundation. Everything that follows, every model, every prediction, every insight, is only as good as this data.

Choose your time buckets

Cash flow forecasting is never one-size-fits-all when it comes to granularity. For daily liquidity management, you want to see flows day by day. For medium-term planning, a weekly view is often more useful. For board reporting and long-range outlooks, monthly or quarterly buckets make the picture easier to read. A good forecasting system should let you switch between these views at will. The practical implication is that data should be stored at the most granular level available (daily), and aggregated up into whichever bucket the user asks for. Once you flatten the data at source, you cannot go back.

Bring In What You Already Know

Your bank data tells you what happened. Your business systems tell you what is going to happen.

Think about all the cash flows in your business that are already committed or highly likely. Open invoices from your ERP system. Scheduled payroll runs. Known tax obligations. Lease payments. Loan repayments. These are cash flows where the amount and timing are already largely determined.

This data lives in your existing systems: your ERP, your HR platform, your treasury management system (which may or may not be Excel). The challenge is getting it out and into your forecast in a structured way.

A good forecasting system connects to your ERP, HR platform, or TMS via API or SFTP to pull this data in automatically. NineAnts does this out of the box, but you can also build these integrations yourself or import CSV files and Excel exports manually. The method matters less than the principle: use what you already know.

The ERP data quality challenge

ERP data is valuable, but it is not always as clean as it looks.

Invoice due dates are not payment dates. Your ERP says an invoice is due in 30 days. But this particular customer historically pays in 45. Which date do you use for your forecast? The contractual due date gives you a best-case scenario. The historical payment behaviour gives you a realistic one. A well-designed forecasting system learns this difference over time. The AI observes that Customer X consistently pays 15 days late and adjusts the timing accordingly.

Partial payments are messy. A customer owes 100,000 but pays 60,000 now and the rest "later." Your ERP might still show the full amount as open. Your forecast needs to account for the split.

Credit notes and disputes muddy the picture. An open invoice of 50,000 means nothing if there is a 20,000 credit note pending against it. Your forecast should work with net amounts, not gross.

Different ERP systems model this data differently. SAP structures payment terms one way. Oracle another. Microsoft Dynamics yet another. If you operate multiple ERPs across your group (which most multinationals do), normalising this data into a consistent format is a challenge.

These are your known cash flows. They form the baseline of your forecast. For the near term, the next few days or weeks, they are usually quite accurate. But the further you look into the future, the thinner this data gets. Open invoices only cover what has already been billed. Payroll only covers what has been scheduled. Three months out, your known cash flows might cover 30% of what will happen.

Known cash flows alone are not going to get you there. That is where AI comes in.

“

I’m a bit obsessed with building better cash flow forecasting. Why?

Because it reduces waste, and no one likes waste.

Done right, it turns information trapped in systems and in people’s heads into something useful.

Less idle cash, less risk and fewer surprises.

And, with a bit of luck, it even brings a little bit of harmony… and maybe some fun.

PHILIP COSTA HIBBERD | CO-FOUNDER, AUTOMATION BOUTIQUE

PART II: THE AI

Let AI Fill the Gaps

This is where it gets interesting.

You now have two layers of data: categorised historical cash flows (what happened in the past) and known cash flows from connected systems (what you already know about the future). The gap between those known cash flows and reality is where AI earns its keep.

How AI cash flow forecasting should actually work

An AI agent analyses your historical cash flow data, category by category, to identify patterns. Three types of patterns matter most:

Seasonality: a repeating calendar-driven pattern. Are your Customer Receipts consistently higher in November and December? Do tax payments cluster in certain months? Does your business see a dip in supplier payments every August when Europe goes on holiday? Seasonality is one of the strongest and most reliable signals for cash flow forecasting because it is driven by structural factors (fiscal calendars, consumer behaviour, industry cycles) that tend to persist year after year.

Trends: sustained directional movements in your data, separate from seasonal fluctuations. Are your customer receipts growing 10% year over year? Are supplier payments gradually increasing as your business scales? Trends capture the underlying trajectory of each cash flow category. A model that captures seasonality but misses the trend will systematically underforecast a growing business.

Recurring patterns: regularities beyond seasonality and trends. Maybe your largest customer pays on the 15th of every month, like clockwork. Maybe capex spikes every 18 months when equipment needs replacing. Maybe you pay dividends twice a year on fixed dates. These are highly predictable patterns that do not fit neatly into a seasonal or trend model but are crucial for an accurate forecast.

Understanding Business Context

Something separates good AI forecasting from great AI forecasting. Patterns alone are not enough. The system needs to understand the business context behind the numbers.

Take a simple example: Rent. Your company has been paying EUR 50,000 per month in rent for two years. Then suddenly, in March, the amount doubles to EUR 100,000. A pure pattern-matching model sees a spike and might treat it as an anomaly, something to smooth over or ignore. But the reality is different. Your company signed a new lease for additional office space. That EUR 100,000 is the new normal. The model should not revert to EUR 50,000 next month. It should project EUR 100,000 going forward.

Now flip it. The same rent category jumps to EUR 150,000 for one month because of a deposit payment on the new lease. That is genuinely a one-off. The system should recognise it as such and not bake it into the recurring forecast.

The difference between these two scenarios, a structural change versus a one-time event, cannot be determined from the numbers alone. It requires understanding what kind of cash flow you are looking at and what is driving the change.

This is where the intelligence in the system matters. A well-designed agentic system does not just look at the time series. It considers the nature of each cash flow category. Rent behaves differently from Customer Receipts. Tax payments follow a regulatory calendar. Capex is lumpy and project-driven. Salaries are stable but step up when headcount changes. Each category has its own logic, and the forecasting approach should reflect that.

In NineAnts, the Treasury Agent classifies the behaviour of each category (stable, seasonal, trending, or volatile) and adapts its approach accordingly. When a pattern changes, it reasons about whether the change is structural or temporary. For a category like Rent, it recognises that a step-change in a previously stable amount likely represents a new baseline, not noise. For a category like Capex, it understands that a large payment is likely project-driven and should not be projected forward unless there is a pattern of similar payments. And when the signal is genuinely ambiguous, e.g. when the numbers alone cannot tell the story, the agent surfaces the change to the treasurer rather than guessing. AI reasoning plus a human in the loop is what handles the edge cases that neither would catch on its own.

This kind of reasoning is what separates a forecasting tool from a forecasting agent. The tool applies a model. The agent thinks about the data.

Which AI Models Work Best for Cash Flow Forecasting

Many forecasting systems, whether built in-house or bought off the shelf, apply the same model to every cash flow category. But a category like Customer Receipts, with hundreds of transactions per month and clear seasonal patterns, needs a fundamentally different approach than a category like Capex, which might have five transactions per year at irregular intervals.

The field of forecasting models is deep. What actually matters for cash flow forecasting falls into three broad families.

Time-series models

These are statistical models purpose-built for sequential data with temporal patterns.

- Meta's Prophet is particularly well-suited for cash flow forecasting because it was designed to handle multiple seasonalities (weekly, monthly, yearly), trend changepoints (moments where the trajectory shifts), and holiday effects. Exactly the kind of patterns that show up in business cash flows.

- ARIMA (AutoRegressive Integrated Moving Average) is a classic approach, especially effective for stationary data where the statistical properties do not change over time.

- Exponential Smoothing methods like Holt-Winters are deceptively simple but surprisingly effective for stable seasonal patterns. Their simplicity also means they are more interpretable. That matters when a treasurer asks "why does the model think this?"

Machine learning models

Approaches like gradient boosting (XGBoost, LightGBM) take a different angle. Instead of treating the data purely as a time series, they can incorporate external features like payment terms, customer segments, and macroeconomic indicators. This makes them powerful when cash flow patterns depend on factors beyond historical timing. Neural networks can capture complex non-linear relationships that simpler models miss, but they need substantially more data to train reliably and are harder to interpret.

These approaches are becoming increasingly relevant as more data becomes available and models become easier to deploy. For most treasury use cases today, time-series and statistical models still deliver the best balance of accuracy and interpretability, but that might shift in the future.

Simple statistical models

Do not underestimate them. A well-tuned moving average or linear regression can outperform sophisticated models for categories with few transactions or very stable patterns. If your rent is 50,000 per month, every month, you do not need Prophet to predict next month. Matching model complexity to data complexity is a sign of engineering maturity, not a compromise.

The agentic approach

In a traditional forecasting system, you would train a model offline on historical data, deploy it, and use it until someone decides to retrain. This approach has a fundamental weakness: the model is always looking at stale assumptions. The world changes, but the model stays frozen.

The agentic approach is fundamentally different. In NineAnts, the Treasury Agent does not train a model once and reuse it. Every time a forecast is requested, the agent analyses the data fresh. It fetches the latest historical transactions, evaluates the characteristics of each category in real time, and selects the most appropriate model on the fly, much like a mechanic reaching for the right tool depending on the task at hand.

This means the forecast always reflects the most current data. If a category that was stable for two years suddenly becomes volatile, the agent detects this on its next forecast run and switches to a model that handles volatility better. There is no waiting for someone to notice, no manual retraining step, no deployment cycle. The agent adapts continuously.

The agent currently selects from multiple model types based on data characteristics:

- Categories with strong seasonal patterns and sufficient transaction volume get time-series models like Prophet
- Stable, low-volatility categories (like rent or insurance) get recurring pattern detection
- Categories with sparse or schedule-driven payments (like quarterly taxes or annual dividends) get schedule-based models
- Categories with very limited data fall back to moving averages

The user can see which model was selected for each category and why. The model fits the data. Not the other way around.

Human in the Loop

Once the agent produces a forecast, a fundamental question emerges: how much control do you want to keep?

This is one of the most important design decisions in any AI system, and it is especially critical in treasury, where the numbers drive real financial decisions. The concept of "human in the loop" is not just a safety feature. It is a philosophy about how humans and AI should collaborate.

At one end of the spectrum, the AI is purely assistive. It generates predictions, but a human reviews and approves every single entry before it enters the forecast. Nothing happens without explicit human consent. This is the right starting point for any organisation that is new to AI forecasting. It builds trust, lets the team learn how the AI thinks, and ensures no erroneous predictions slip through unnoticed.

In the middle of the spectrum, the AI operates semi-autonomously. It fills in forecasts automatically for categories where it has high confidence, but flags uncertain predictions for human review. The human focuses their attention where it matters most, rather than reviewing every number.

At the other end, the AI operates with high autonomy. It injects predictions into the forecast, handles routine categories end to end, and only escalates to the human when something unusual appears. The human sets the boundaries and monitors the outcomes, but does not approve every individual entry.

In NineAnts, autonomy is not a simple low/high toggle. The Treasury Agent can be configured at a granular level: for each type of action, e.g. generating a forecast, writing a prediction directly into a category, adjusting existing entries, running scheduled jobs, sending notifications, you decide whether the agent may act on its own, must ask for confirmation first, or is not allowed to perform the action at all. This means an organisation can run the agent in a conservative mode for sensitive categories while letting it operate more independently on the routine ones.

In practice, most teams start with the agent in a review-first posture, where every prediction is presented as a suggestion that the user accepts, modifies, or rejects before it enters the forecast. As confidence grows, they unlock specific actions, e.g. letting the agent automatically inject forecasts for stable categories while still requiring approval for volatile ones. Regardless of the settings, every AI-generated entry remains clearly marked, and the user can always review and override after the fact.

Regardless of the autonomy level, every AI-generated forecast entry is clearly marked with its source. The user always knows which numbers came from the AI, which came from connected systems like the ERP, and which were entered manually. This transparency is essential. If you cannot distinguish AI-generated data from human-entered data, you cannot meaningfully review or trust the forecast.

The right autonomy level is not a permanent choice. Most organisations start at low autonomy, build confidence over a few forecast cycles, and gradually increase it as they see the AI perform. The goal is not to remove the human. It is to focus the human on the decisions that need human judgment.

From reactive assistant to proactive team member

There is another dimension of autonomy that goes beyond how the agent handles individual forecasts. It is about when the agent works.

Most AI tools are reactive. You open the application, type a question, and get an answer. The AI waits for you. It does nothing until you ask. This is useful, but it limits the AI to the role of an assistant: someone who helps when called upon.

A truly agentic system can also be proactive. It can run on a schedule, without anyone opening the application. It can generate the weekly forecast every Monday at 7 AM, before the treasury team arrives. It can compare last week's actuals against the forecast every Friday afternoon and flag the biggest deviations. It can monitor bank balances daily and alert the team when excess cash exceeds a threshold. It does not wait to be asked. It does the work and reports back.

This is the difference between an assistant and a team member. The assistant answers questions. The team member takes initiative.

In NineAnts, users can configure scheduled agent tasks using natural language. Tell the agent to "create a forecast for all categories every Monday at 9 AM" and it will. The agent runs autonomously on that schedule, executes the forecast, and can email the results directly to the treasury team or any stakeholder. No one needs to be logged in. No one needs to click a button. The forecast is ready when you need it.

This turns the agent from something you interact with into something that works alongside you. It handles the routine, recurring work autonomously, and frees the team to focus on the exceptions, the judgment calls, and the strategic decisions that actually need a human.

The User Experience

Something that does not get enough attention in conversations about AI forecasting: the user experience.

A cash flow forecast is only useful if the people working with it trust it. And trust comes from clarity. When a treasurer opens the forecast on Monday morning, they are looking at numbers that come from fundamentally different sources. Some entries were pulled automatically from the ERP. Some came from bank statement data. Some were manually entered by a colleague in the Singapore office. Some were generated by an AI model. And some are a combination, where an AI prediction was manually adjusted by a local finance team.

If all of these numbers look the same on screen, you have a problem. The treasurer does not know what to question, what to verify, and what to trust. A manually entered number from a subsidiary has a very different confidence profile than an AI prediction based on 24 months of seasonal data. An ERP figure based on booked invoices is more certain than a forecast for a category where no system data exists. Treating them all the same is misleading.

This is a design problem, not a technology problem. And solving it well is one of the most important things you can do when building a forecasting system.

Every number needs a label

At a minimum, every forecast entry should clearly show its source. Was it imported from the ERP? Pulled from a bank feed? Entered manually by a user? Generated by the AI? Uploaded from a file? The user should never have to guess.

In NineAnts, every forecast entry carries a source type (AI, ERP, bank, manual, file upload) that tells you exactly where it came from. The interface makes this visible at a glance. AI-generated entries look different from system-imported entries, which look different from manual entries. You can filter by source, so if you want to review only the AI-generated numbers, you can isolate them instantly.

Audit trails matter more than you think

Source labels tell you where a number came from. But in a collaborative forecasting process, you also need to know who changed what and when. Did someone in the French subsidiary override an AI forecast last Tuesday? Did the treasury team adjust the investment assumptions after the board meeting? Was a category re-forecasted because new ERP data came in?

A full audit trail turns the forecast from a static spreadsheet into a living, traceable document. Every action (create, update, delete, import, AI injection) should be logged with a timestamp and a user. This is not just good practice. For many regulated industries, it is a compliance requirement.

Designing for collaboration across subsidiaries

The user experience challenge multiplies in a decentralised setup. When twenty subsidiaries contribute to a single consolidated forecast, each user needs to see their own scope clearly. The finance manager in Brazil should see their entities and their categories, not the entire group forecast. But group treasury needs the consolidated view with the ability to drill down into any entity.

The interface needs to serve both audiences without overwhelming either. Local users need simplicity: their categories, their numbers, what the AI suggests, approve or adjust. Group treasury needs depth: consolidated totals, drill-down by entity, filter by source type, and the ability to spot which subsidiaries have submitted and which have not.

Getting this right is what separates a tool that people actually use from one that gets abandoned after the first quarter. The best forecasting model in the world is worthless if the people who need to interact with it cannot understand what they are looking at.

Control matters just as much as collaboration. Central treasury should be able to close a forecast cycle and lock it against further edits once the submission window ends. Coherence of data across entities is more important than squeezing in a last-minute change from one subsidiary. A partial late update (for example, one entity adding intercompany flows while its counterpart does not) creates mismatches that are worse than having no update at all. Clear ownership of what is expected, by whom, by when, and under which guidelines is what makes a decentralised process work in practice.

The Role of the LLM in an Agentic Forecasting System

The word "agent" has been used throughout this article. It is not marketing language. It is a specific technical architecture that fundamentally changes how forecasting systems work.

Traditional forecasting software is procedural. You configure it. You press a button. It runs a model. It gives you output. If the output does not make sense, you adjust the configuration and try again. The software does exactly what you tell it to do, nothing more.

An agentic system is different. At its core sits a Large Language Model (LLM) that acts as the brain of the system. The LLM does not do the statistical forecasting itself. Prophet (an open-source time-series forecasting library originally developed by Meta) handles the time-series modelling. Moving averages handle the sparse categories. The forecasting models are specialised tools. But the LLM is the orchestrator. It decides which tool to use, when, and why.

What the LLM actually does

When you ask the system to forecast Customer Receipts, something happens behind the scenes. The LLM first classifies your intent. Are you asking for a forecast, an analysis, a comparison, or something else? Then it retrieves the relevant data: historical transactions, existing forecasts, connected source data. It examines the data characteristics and selects the appropriate forecasting model. It interprets the model's output in the context of what it knows about the category. And it communicates the result back to you in plain language, explaining not just the number but the reasoning behind it.

This is fundamentally different from a button that runs Prophet and shows you a chart.

The LLM can also reason about things that statistical models cannot. Is this category seasonal or volatile? Does the recent data suggest a trend change? Is there a discrepancy between the forecast and the actuals that deserves attention? These are judgment calls that require understanding the meaning of the data, not just its mathematical properties. An LLM can read a transaction description, understand that "Office Lease Q2" is rent-related, and categorise it accordingly. A pure rule engine would need a specific rule for that.

AI security and data governance

Every treasury team considering AI should be asking hard questions about security. And they should be sceptical of any vendor that brushes those questions aside.

Financial data is among the most sensitive information a company holds. Transaction histories, cash positions, bank account details, counterparty relationships, forecast assumptions. This is the kind of data that, if mishandled, creates real risk. Not theoretical risk. Career-ending, board-level risk.

Traditional treasury software already raised security concerns. AI raises new ones. The data is not just stored and displayed. It is processed, analysed, and used to generate outputs that did not exist before. That changes the threat surface. And it means the questions you need to ask go beyond the standard IT security checklist.

Here are the concerns we hear most often from corporate treasury teams, and what you should be thinking about whether you are building an AI forecasting system in-house or evaluating an external solution.

Where does my data go?

The first question is deceptively simple. Where is my data hosted? Which data centres? Which jurisdiction? Who has access?

Data residency requirements vary by region. In Europe, GDPR and internal policies often mandate that data stays within the EEA. In the US, industry regulations and state-level privacy laws impose their own constraints. Across Asia-Pacific, rules differ country by country. Whatever your jurisdiction, the principle is the same: you need to know exactly where your data is processed, not just stored. With AI, the question becomes more nuanced. Your application data might sit in Frankfurt or Virginia, but if the AI model runs on servers in a different jurisdiction, your prompts, inputs, and outputs may be crossing borders every time a forecast is generated.

What to look for: a clear statement that all processing, including AI inference, happens within infrastructure you can verify. Not just data at rest, but data in motion. Every API call, every model interaction, every generated output.

Is my data used to train AI models?

This is the question that keeps Chief Information Security Officers (CISOs) up at night. And for good reason.

Many AI providers use customer interactions to improve their models. That means your financial data, your transaction patterns, your forecast logic could end up embedded in a model that serves other customers. Even if it is anonymised or aggregated, the risk is real and the optics are worse.

Any serious vendor should give you an unambiguous contractual guarantee: your data is not used for model training, fine-tuning, or service improvement in ways that could expose it to others. No exceptions, no grey areas.

There is a meaningful distinction here between using your data to train a foundation model (which should never happen) and using interaction signals like acceptance or rejection of a forecast to improve workflow logic for your own tenant. The second is how the system learns your preferences. The first is a red line.

Can other customers see my data?

In a multi-tenant SaaS platform, multiple customers share the same infrastructure. That is normal and cost-effective. But when AI is involved, the isolation question gets more complex.

A traditional application serves data from a database. The query either returns your data or it does not. But an AI model generates outputs dynamically. If the model has been exposed to data from multiple tenants during processing, there is a risk of cross-contamination. An AI output could, in theory, contain fragments or patterns derived from another customer's data.

What to look for: logical tenant separation at every layer, including the AI layer. Input validation and output filtering to prevent prompt injection attacks. Clear safeguards against cross-tenant data exposure in AI-generated outputs. And incident management procedures that specifically address AI-related risks, not just traditional data breaches.

How much can the AI do without my approval?

We covered the topic of human oversight and autonomy levels in detail in the section “Human in the Loop.” From a security perspective, the key point is this: the more autonomy an AI system has, the more important it is that you can control and audit what it does.

In treasury, the numbers drive real decisions. A forecast that is automatically pushed into a board report, an investment recommendation that triggers an execution, a categorisation change that shifts your cash position by millions. These are not areas where you want “the AI just did it” as an explanation.

Under regulations like GDPR in Europe, CCPA in California, or sector-specific rules elsewhere, automated decisions with significant effects may require specific safeguards, including the right to human intervention. Even if your use case does not hit a legal threshold, the principle is sound. You should be able to configure exactly what the AI can do on its own, what requires approval, and what is off limits entirely.

Can I trace what the AI did and why?

The section “The User Experience” already explains why audit trails and source labels matter for building trust in your forecast. From a security and compliance angle, AI adds a specific requirement on top: you need to know not just that a number changed, but which AI model generated it and what data it was based on.

What to look for: logging of AI-generated outputs that includes the model or workflow used, the data sources referenced, and any subsequent human modifications. This is especially important in regulated industries, but even without a compliance mandate, it is what allows you to defend any number in a board meeting.

The questions that matter

Whether you are evaluating a vendor or building internally, these are the questions your security and compliance teams should be able to answer. If you are buying, ask the vendor. If you are building, ask yourself.

Questions to ask

1. ***Where exactly is my data hosted, including AI processing? Which data centres, which jurisdiction?***
2. ***Is my data ever used to train, fine-tune, or improve AI models?***
3. ***How is my data separated from other customers, specifically in the AI layer?***
4. ***Can I control what the AI does autonomously versus what requires my approval?***
5. ***How are AI-generated outputs logged and for how long?***
6. ***What happens to my data when I terminate the contract?***
7. ***Do you have an ISO 27001 certification or equivalent?***
8. ***How do you handle AI-specific incidents like cross-tenant data exposure or prompt injection?***
9. ***Can I bring my own AI model, host it on my own infrastructure, or swap providers later?***
10. ***If buying: will the vendor share their Data Processing Agreement, including any AI-specific annexes, before you sign? If building: have you documented equivalent controls internally?***

If you are buying and the vendor cannot answer these clearly, that tells you something. If you are building and your own team cannot answer them, you have work to do before going live.

How we handle this in NineAnts

We built NineAnts with the assumption that every customer would ask these questions, and that we should be able to answer all of them before the first meeting ends.

All data is hosted on Microsoft Azure infrastructure, with data residency configured to match customer requirements. For European customers, this means EEA-based data centres. For customers in other regions, we work with you to ensure data stays where it needs to be. AI processing, including all LLM interactions, runs within infrastructure that Automation Boutique controls. Customer data is never sent to third-party AI providers as independent recipients or controllers. It is not used to

train, fine-tune, or improve AI models. Tenants are logically segregated at every layer, including input validation and output filtering in the AI pipeline.

As described in the section on human-in-the-loop, autonomy in NineAnts is configurable per action type. From a security standpoint, this means customers control exactly what the AI is allowed to do, and every AI-generated output is logged with full traceability.

Automation Boutique maintains an ISO 27001 certified Information Security Management System. Our Data Processing Agreement includes a dedicated annex on AI security and processing controls, covering everything from model training restrictions and tenant isolation to AI-specific incident management and DPIA support. We share this documentation with every customer during onboarding, because we believe transparency is not a nice-to-have. It is the baseline.

The Agent Architecture

Think of it as a team. The LLM is the team lead. The forecasting models are the specialists. The data connectors are the researchers. The action engine is the analyst who translates findings into recommendations. The team lead coordinates all of them, routes the right questions to the right specialist, synthesises the answers, and communicates with the user.

In practice, this means the system has dozens of specialised nodes, each handling a specific part of the workflow, with the LLM routing between them based on what the situation requires. Ask for a forecast, and you get one path. Ask why last month's actuals deviated from the forecast, and you get a completely different path through the system. The LLM figures out the right path for every query.

This architecture is what enables the system to handle ambiguity, unexpected questions, and workflows that a traditional software designer could not anticipate. A treasurer might ask: "Why were our supplier payments so high in March?" That is not a forecast request. That is an analytical query that requires pulling actuals, comparing to historical norms, identifying the outlier transactions, and forming an explanation. A procedural system cannot handle this. An agentic system can.

This is the revolution in treasury technology. Not just better models. Better reasoning.

The agent architecture

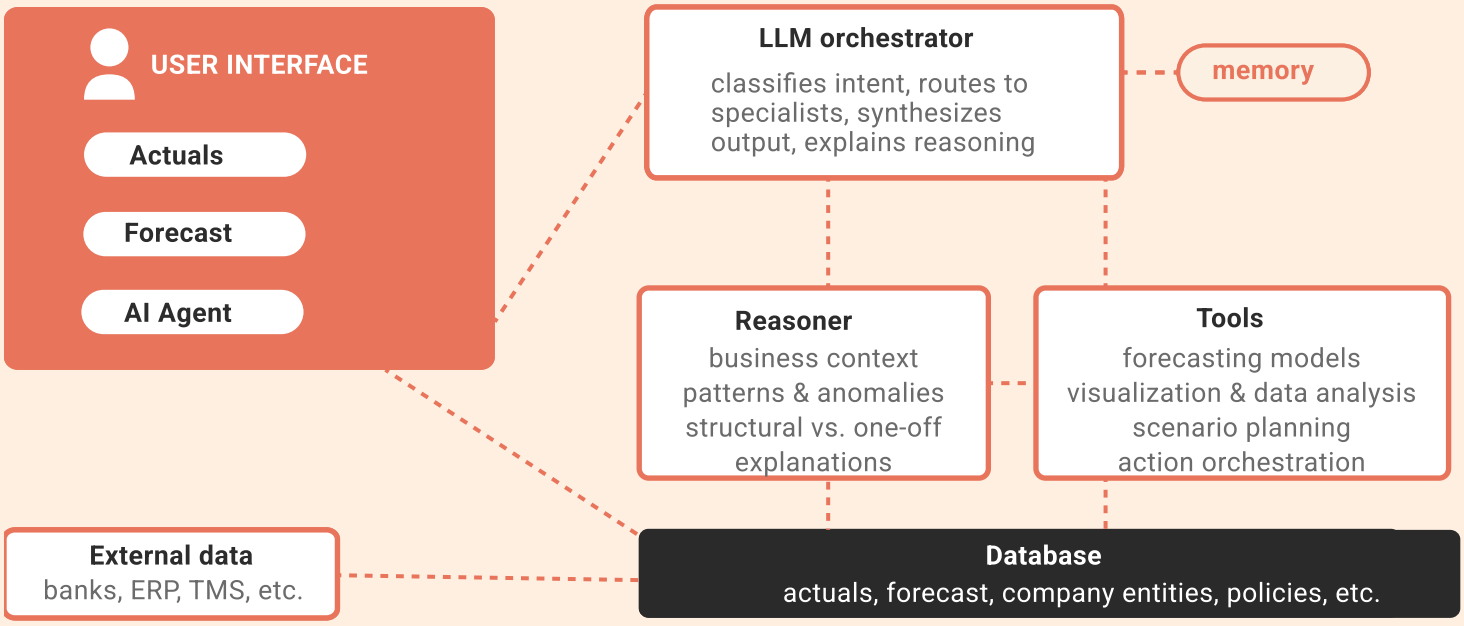

PART III: THE ORGANISATION

Multi-Entity and Multi-Currency Forecasting

Forecasting for a single entity in a single currency is one challenge. Consolidating forecasts across ten, twenty, or fifty entities in multiple currencies is a different game entirely.

This is where many forecasting solutions, and many Excel-based processes, start to break down. The complexity does not scale linearly. It multiplies. Three challenges in particular are worth calling out: eliminating intercompany flows, handling multiple currencies, and reconciling timing mismatches between entities.

Intercompany eliminations

Entity A in the Netherlands pays 500,000 to Entity B in Germany for management services. That is an outflow for A and an inflow for B. But at group level, it nets to zero. No cash enters or leaves the group.

If your consolidated forecast does not eliminate intercompany flows, you are double-counting. Your cash inflows look 500,000 higher than reality. Your cash outflows look 500,000 higher than reality. The net effect might wash out on a total cash flow number, but the individual category totals are inflated, and any analysis built on those category totals is wrong.

The cleanest way to handle this is not to try to match and eliminate every intercompany flow after the fact, but to tag them up front. Every transaction between group entities goes into a dedicated Intercompany category. At a single-entity view, those flows still appear as normal inflows and outflows. At the consolidated group level, the Intercompany category should net to zero. If it does not, that is your signal to drill down and find the missing or mismatched leg. This keeps the responsibility for correct tagging with the entities that booked the transaction, rather than asking the consolidation layer to guess which outflow matches which inflow.

Multi-currency challenges

Each entity forecasts in its local currency. Your subsidiary in Sweden forecasts in SEK. Your operation in the US forecasts in USD. Your Brazilian entity forecasts in BRL. Group treasury needs the consolidated view in the reporting currency, typically EUR or USD.

Which exchange rate do you use to convert? This is not an academic question. The answer materially changes your forecast.

- **Spot rates** give you today's exchange rate. Simple, but your forecast for three months out will change every day as rates move, even if the underlying business forecast has not changed at all.
- **Budget rates** are the exchange rates set at the beginning of the fiscal year for planning purposes. They provide stability because your forecast does not swing with daily FX movements. But they may diverge significantly from reality as the year progresses.
- **Forward rates** represent the market's expected exchange rate for a future date. They are the most economically accurate for forecasting purposes, but they require a data source for forward rate curves and add complexity to the system. For actuals rather than forecasts, you face a similar choice: you can translate every transaction at the spot rate of the day it settled, or apply a single closing rate to the whole period. Both are valid, but they answer different questions about FX performance.

Consider a Nordic group with entities in SEK, NOK, DKK, and EUR. Each entity submits their forecast in local currency. Treasury needs the EUR consolidated view. The AI converts at the applicable rates, and flags entities where FX movements have materially changed the EUR-equivalent forecast since the last submission. What used to take a full day of manual work in Excel now happens in minutes.

Timing mismatches

There is one more subtlety that trips people up. Entity A books a payment on Monday. Entity B receives it on Wednesday. In between, there is a two-day window where the cash has left one account but not arrived at the other. In the consolidated forecast, this creates a temporary phantom cash gap that does not actually exist in reality.

For intraday or daily forecasting, these timing mismatches can be significant. Your system needs to understand that a payment and its corresponding receipt are the same transaction at different stages of settlement. Getting this right is the difference between a consolidated forecast that shows reality and one that shows artifacts.

Centralised vs. Decentralised Forecasting

Who submits forecast data in your company? Does treasury build the entire forecast centrally? Or do subsidiaries, business units, and divisions each submit their piece?

There is no universally right answer, but from experience, bottom-up usually wins. Local teams have better visibility on local conditions. A subsidiary in Brazil knows more about their supplier payment cycles than headquarters in Amsterdam. A business unit running a major project knows more about the timing of their capex than a centralised treasury team. The people closest to the business are usually the best source of forecast data.

The hybrid approach

The most effective model is a hybrid: central treasury sets the framework (the categories, the time horizons, the submission deadlines, the quality standards) and local teams provide the inputs within that framework. The AI validates everything.

This gives you the accuracy benefits of local knowledge with the consistency benefits of a centralised structure. The framework ensures everyone is speaking the same language. The local inputs ensure the data reflects reality.

AI as the quality layer

Decentralised forecasting introduces a quality problem. If twenty subsidiaries each submit forecasts, how do you ensure consistency and accuracy across all of them?

This is where AI becomes incredibly valuable. The treasury agent can automatically review every submitted forecast against historical patterns and flag anomalies.

A concrete example. A subsidiary in Brazil submits their quarterly forecast. The agent notices that supplier payments are 40% below the same period last year, with no explanation and no corresponding change in customer receipts. It flags this. The treasury team reaches out. It turns out the local team forgot to include a large vendor contract renewal that is due next month.

In NineAnts, the Treasury Agent can review submitted forecast data and surfaces likely inconsistencies. Did a subsidiary forget to include their quarterly tax payment? Are the numbers wildly different from prior periods? Has a recurring pattern suddenly disappeared? The agent prompts subsidiaries to double-check their numbers when something looks off.

These guardrails do more than catch errors. They educate users across the organisation about what a good forecast looks like. Over time, the quality of submitted forecasts improves because people learn what the agent checks for. The subsidiary that forgot the vendor contract renewal will not forget it next quarter. This is behaviour change at scale, driven by AI but ultimately making the humans better at their jobs.

Common Pitfalls and How to Avoid Them

We have seen many companies implement AI cash flow forecasting. The ones that succeed avoid these pitfalls. The ones that struggle usually fall into at least two of them.

Pitfall 1: Garbage in, garbage out

This is the most common failure and the most preventable. Companies get excited about AI and want to start forecasting immediately. Before their historical data is clean. Before their transactions are properly categorised. Before they have enough history to detect patterns. The AI model produces outputs, so it looks like it is working. But the outputs are built on bad data, which means the forecasts are unreliable, which means the team loses trust, which means the project gets shelved.

How to avoid it: spend the time on the foundation. Get 24 months of clean, categorised history before you turn on forecasting. It is not glamorous, but it is the highest-leverage thing you can do.

Pitfall 2: Over-trusting the model

AI forecasts are not guarantees. Usually a forecast of 5 Million next week will be more accurate than 5 Million in one year. In both cases, there is no promise that exactly 5 Million will arrive. It means the model's best estimate is 5 Million, with some range of uncertainty around it.

The difference between a very tight estimate and a very wide one is enormous. The first one you can confidently invest against. The second one you should treat as directional guidance at best. Always understand how confident the system is in its prediction, and treat different categories differently. Some categories forecast beautifully. Others are inherently noisy.

Pitfall 3: Ignoring business context

AI is extraordinarily good at finding patterns in historical data. A well-designed agent can pick up signals that pure statistical models would miss.

That said, the AI does not know about the acquisition you are planning next quarter, or the customer you are about to lose. These are contextual inputs that require human judgment. The best AI forecasting systems make it easy for users to overlay their knowledge on top of the AI's predictions. The AI provides the baseline. The human provides the context. Together, you get a better forecast than either could produce alone.

Pitfall 4: Trying to forecast everything at once

It is tempting to roll out AI forecasting across every category, every entity, and every time horizon simultaneously. A better approach is to start with rough, broad categories that cover the bulk of your cash flows and refine them as you learn what the data looks like. Alternatively, focus first on the business units or legal entities that drive 80% of your cash flow, get the forecast working there, and then expand to the rest of the group. Prove the value, build trust, and grow the scope deliberately.

Customer Receipts and Supplier Payments alone typically represent 60-70% of total cash flows. Start there. Add Salaries and Taxes next. Then the long tail. This phased approach lets you deliver results quickly while managing complexity.

Pitfall 5: No feedback loop

If you never compare your forecast to what actually happened, you never improve. And worse, you never know how wrong you are. We have seen companies produce forecasts for years without ever systematically measuring forecast accuracy. They believe their forecast is "pretty good" because they have never measured how bad it actually is.

How to avoid it: close the loop. Every period, compare actuals to forecast. Measure accuracy by category. Identify systematic biases. This is the flywheel, and it is the single most important practice for long-term forecast quality.

This is why we built automated feedback loops into NineAnts from day one. But even without a specialised platform, you can avoid these pitfalls with discipline and the right processes.

PART IV: THE VALUE

From Forecast to Action

A cash flow forecast is not the finish line. It is the starting line.

The entire purpose of forecasting is to identify risks and opportunities early enough to act on them. This is where the real business value lives. And this is where most companies under-invest. They spend all their energy building the forecast and then struggle to extract actionable insights from it.

Excess cash

If your forecast shows surplus liquidity for the next three months, that cash should not be sitting idle in a zero-interest account. It should be working. Money market funds, term deposits, or paying down a revolving credit facility. The right instrument depends on the amount, the duration, and your company's investment policy.

An intelligent agent can take this further. The forecast shows 8 Million of excess cash for the next sixty days. The agent analyses current money market rates across your approved counterparties, checks your investment policy limits (maximum 30% with any single counterparty, minimum credit rating of A, maximum tenor of 90 days) and recommends splitting: 5 Million in a 30-day term deposit at 3.2% with Bank A, and 3 Million in an overnight deposit as a liquidity buffer. The treasurer reviews, approves, and executes. What used to require hours of analysis and rate comparison happens in seconds.

Cash shortages

On the other side, if the forecast shows a funding gap in eight weeks, you have time to arrange a credit facility, draw on a revolving credit line, accelerate receivables or delay payables. Eight weeks of lead time is the difference between a planned financing decision and a panic call to the bank.

FX exposure

For companies operating in multiple currencies, the cash flow forecast reveals currency mismatches. If you are forecasting EUR expenses and USD receipts, you have FX risk. Identifying it in the forecast means you can hedge it before it becomes an FX loss on your P&L. The forecast tells you the notional amount and the timing. The agent can recommend specific hedging instruments aligned with your hedging policy.

Early warning system

Perhaps most importantly, a cash flow forecast is an early indicator of business health. If a division's cash flows are deteriorating, customer receipts dropping and supplier payments stretching, that is a signal. Management can act before the situation becomes a crisis. This is the kind of insight that turns treasury from a back-office function into a strategic partner to the business.

Our Treasury Agent at NineAnts goes beyond simply flagging "you might have excess cash." It analyses the size, duration, and investable amount, then translates that into concrete investment scenarios. It calculates FX hedging recommendations based on your forecasted exposures and your company's hedging policy. It identifies negative trends in specific business units before they become visible in the P&L. Actionable recommendations, not just observations.

The Flywheel: Actuals vs. Forecast

Most people miss this. AI cash flow forecasting is not a one-time exercise. It is a flywheel. And the flywheel is what separates companies that get marginal value from AI and companies that get transformational value.

Every day, new actuals come in. Real transactions hitting your bank accounts. These actuals can be automatically compared against what was forecasted. Where was the forecast accurate? Where did it miss? By how much? And why?

This is called variance analysis, and it is traditionally one of the most time-consuming tasks in treasury. Comparing actuals to forecast across dozens of categories, multiple entities, and various time horizons requires enormous manual effort. Most treasury teams do it monthly, if they do it at all.

An AI agent does this continuously. It detects accuracy gaps and significant deviations. A large capital expenditure that was never forecasted. A customer receipt that came in two weeks late. It identifies patterns in the forecast errors themselves. Are you consistently overforecasting supplier payments? Underforecasting tax obligations?

These meta-patterns are surfaced to the treasurer as clear feedback to act on — adjusting assumptions, refining categories, or correcting specific source data — so that the next forecast cycle starts from a better place. (The model itself is not being retrained on your data; the improvement comes from the humans getting better, helped by the agent.)

The compounding effect is powerful. Each cycle of forecast, actuals, analysis, and learning makes the next forecast better.

Forecast Versions

A cash flow forecast is not a single, static document. It is a living process that evolves over time. Every week, new data comes in. Assumptions change. Subsidiaries update their inputs. The AI refines its predictions. Each of these moments creates a new version of the forecast.

Tracking different versions of your forecast is essential for two reasons.

First, accountability. When the CFO asks "how has our outlook changed since last month?", you need to be able to pull up both versions and show exactly what changed, where, and why. Without version history, every forecast exists in isolation, and you lose the ability to track how your understanding of the future evolves.

Second, accuracy measurement. You can only measure how good your forecasting process is if you can compare what you predicted at different points in time against what actually happened. A forecast made six weeks ago should be less accurate than a forecast made two weeks ago for the same period. If it is not, something is wrong with your process.

How often should you create a new forecast version?

The cadence of your forecast versions depends on your business, your reporting cycles, and how dynamic your cash flows are. There is no single right answer, but there are clear patterns.

Many treasury teams create a new forecast version weekly. This works well for companies with active, transaction-heavy businesses where the cash flow picture changes meaningfully from week to week. Every Monday, a new version is created. The latest data from connected systems is pulled in. The AI fills the gaps for the forecast horizon. The team reviews and adjusts. By Tuesday morning, the organisation has a fresh, up-to-date view of the coming weeks and months.

Other organisations forecast monthly, aligned with their financial reporting cycle. A new version is created at the start of each month, incorporating the previous month's actuals and refreshing the forward-looking predictions. This is common in companies with more stable, predictable cash flows or in organisations where treasury operates on a monthly rhythm.

Some companies create forecast versions tied to specific events: before a board meeting, before a funding decision, before a major payment run. These event-driven versions capture a point-in-time view that can be referenced later.

The key principle is that each version is a snapshot. Once created, it can be locked so that no further changes are made. This preserves the integrity of each version for comparison purposes. You want to be able to look back and say: this is exactly what we forecasted on March 1st. If versions keep getting edited after creation, you lose the ability to track how your forecasting accuracy evolves.

In NineAnts, creating a new forecast version takes seconds. The agent can do it on a schedule, automatically pulling in the latest data from connected sources and generating AI forecasts for each category. Previous version data can optionally be carried forward as a starting point, so the team is not starting from scratch each time. Versions can be locked once finalised, and the agent tracks the full history for comparison.

Forecast vs. forecast analysis

Comparing different forecast versions against each other reveals something that actuals vs. forecast analysis cannot: the stability and reliability of your forecasting process itself.

Are your forecasts converging as the period approaches? That is a sign of a healthy process where each iteration adds information. Are they volatile right up until the last moment? That suggests your inputs are unreliable or your process is not capturing information early enough.

This analysis also reveals which categories are forecastable and which are not. If Customer Receipts are consistently stable across versions but Capex swings wildly, that tells you where to focus your improvement efforts. It might also tell you that Capex needs a fundamentally different forecasting approach, or that the business needs to provide better visibility into planned capital expenditures.

In NineAnts, every forecast is stored as a versioned snapshot. The agent can compare any two versions side by side, highlighting the differences by category and period. It tracks how forecasts evolve over time and surfaces insights about forecast stability that would take hours to produce manually.

Scenario Analysis: What If?

A cash flow forecast tells you what is most likely to happen. But treasurers do not just need the base case. They need to understand what happens if things go differently than planned.

What if customer receipts drop by 20% next quarter? What if a major supplier changes payment terms from 30 to 60 days? What if the EUR/USD rate moves 10% against you? These are the questions that keep treasurers up at night, and they are exactly the questions that scenario analysis is designed to answer.

How agents make scenario analysis practical

Traditionally, scenario analysis in treasury meant building three versions of the forecast in Excel: best case, worst case, and base case. Each version required manual adjustments across every category, every entity, and every time period. By the time you finished the worst case, the base case was already outdated.

An agentic system makes this easier. You describe the scenario in plain language or adjust a few parameters, and the agent recalculates the entire forecast accordingly. Want to see what happens if Customer Receipts decline by 15%? The agent applies the adjustment across all relevant periods, recalculates your cash position, and immediately shows you the impact on your liquidity, funding needs, and investment capacity.

The most valuable scenarios are not just "what if revenue drops by X%." They are compound scenarios that reflect real business situations:

- *What if we lose our largest customer AND the EUR weakens by 5%?*
- *What if supplier payments accelerate because of tighter credit terms across the industry?*
- *What if we delay the planned acquisition by one quarter?*

Each of these scenarios has cascading effects across the forecast. Losing a customer affects receipts but also changes the FX exposure profile. Accelerating supplier payments affects cash outflows but also changes the investment window. An agentic system can model these compound effects, giving treasury a much richer picture of their risk landscape.

In NineAnts, the agent supports scenario modelling where you can adjust forecasts by category or across the board, and instantly compare the scenario against the baseline. You see both the adjusted forecast and the original side by side, making it easy to quantify the impact of each scenario.

PART V: THE FUTURE

Where Do We Go From Here?

Everything in this playbook, the data foundation, the connected sources, the AI forecasting, the multi-entity consolidation, the quality checks, the variance analysis, builds toward something bigger.

If your forecast is accurate, and your agent identifies the optimal actions, the natural next question is: why not execute them?

Think about it. The agent identifies 5 Million of excess cash for the next 45 days. It knows your investment policy. It knows the current money market rates. It knows your counterparty limits. It recommends a specific term deposit with a specific bank at a specific rate. Today, a treasurer reviews that recommendation and manually executes the trade.

But what if the agent could do it? What if, within the boundaries you set, it could automatically execute the investment? Or place the FX hedge? Or draw on the credit facility?

This is not science fiction. This is the next frontier of treasury automation. And the cash flow forecast is the foundation that makes it possible.

The implications extend beyond individual transactions. Treasury becomes a real-time, continuously optimising function. Not a team that produces quarterly reports and reacts to surprises, but an intelligent system that anticipates, recommends, and, within defined boundaries, acts. Investment portfolios are continuously rebalanced as the forecast evolves. FX hedges are adjusted as exposures change. Funding is arranged before shortfalls materialise. The treasury team shifts from executing routine decisions to setting strategy and handling exceptions.

Remember the CFO who walked into the quarterly business review and asked: "Where will our cash be in thirteen weeks?" Back then, the answer took two weeks of spreadsheet consolidation and was already outdated when it arrived. Now imagine that same question. The agent already has the answer. Updated this morning. Broken down by entity, by currency, by category. With a confidence range and a list of recommended actions attached. The CFO does not wait. The treasury team does not scramble. The answer is already there.

That is the future this playbook is building toward. Not a single tool or a single model, but a system that turns financial data into decisions, and decisions into action. Step by step, with humans setting the boundaries and AI doing the heavy lifting within them.

About us

Automation Boutique builds AI solutions for treasury and finance teams.
Our agentic platform, NineAnts, is built around AI cash flow forecasting, the topic of this guide, and extends naturally into more autonomous treasury operations. It is designed from day one around clean data, explainable reasoning, and human-in-the-loop mechanisms.

Enterprise-grade security is built in. Automation Boutique is ISO 27001 certified, and the same controls extend to every AI model, agent and data flow inside NineAnts.

Learn more at nineants.com or automationboutique.com.