

The 2026 Mid-Market Planning Landscape: *a practitioner's view*

What the published evidence says about how finance teams plan in 2026 — spreadsheet dependency, forecast accuracy, the close, the distance between AI intent and AI in use, and what the tooling market actually costs and takes to deploy — read through the experience of people who have implemented planning systems at enterprise and mid-market scale. Every figure is cited to a named source with its date and sample; where a source covers large enterprises, we say so and say what we think carries down to a company of 100 to 1,500 people.

AUDIENCE	METHOD	SOURCES	READING TIME
Leaders of mid-size enterprises	Synthesis of published surveys, 2025–2026	16 named, with sample sizes where published	16 min read

Series Hours, Not Months · Flagship

Standard no primary survey of our own; third-party figures cited by source; enterprise-sample findings flagged; opinion labelled as such

00 Executive summary

Five findings, each carried by more than one source, and each with a consequence for a mid-size company deciding what to do about planning this year.

FINDING 1

The spreadsheet is still the planning system

96% of FP&A professionals use spreadsheets as a planning tool at least weekly and 93% for reporting ^[1]. This is not a legacy tail; it is the mainstream, and every planning platform is adopted alongside it, not instead of it.

FINDING 2

Forecast accuracy is the cost problem CFOs name first

61% of CFOs say inaccurate forecasting is the single biggest barrier to controlling costs, and 82% of companies report making decisions on stale information ^[2]. The decade-old finding that finance spends roughly three-quarters of its time gathering and validating data rather than analysing it ^[3] has not been overturned.

FINDING 3

AI intent is near-universal; AI in core workflows is not

87% of large-company CFOs expect AI to be very or extremely important to finance in 2026 and 54% make integrating AI agents a top transformation priority ^[4] — while only 17% of finance teams use AI in core workflows and 45% remain in limited pilots ^[5], and 23% use AI in forecasting regularly ^[1].

FINDING 4

The mid-market tooling market is priced and paced for the mid-market's larger cousins

Median annual contracts for the mid-market planning tools run roughly \$33,000–\$42,000, with implementation from nothing to \$40,000 and more ^[6]; deployment timelines of two to four months for mid-market tools and six to twelve for enterprise platforms are the published norm ^[7]. 69% of finance transformation programmes are reported behind schedule ^[2].

FINDING 5

The teams pulling ahead did three unglamorous things

Across the surveys' own descriptions of leading teams: they centralised planning data, automated the manual chain between ledger and plan, and applied AI to specific forecasting problems rather than general experimentation ^[2]^[5]. None of the three requires a large company.

CONSEQUENCE

The mid-market can now do what only the enterprise could afford

Our opinion, developed in sections 05 to 07: the constraint on a 100–1,500-employee company is no longer capability or price but implementation time and governance — and both are now tractable in weeks rather than quarters if the platform is one system rather than three.

01 The state of practice

02 The forecasting gap

03 The close and the data chain

04 AI: intent, adoption, use

05 The tooling market

06 What the leaders do differently

07 A readiness self-assessment

08 What this means for a mid-size company

09 Method and sources

A note on samples

The strongest CFO surveys skew large. Deloitte's CFO Signals polls 200 finance chiefs at North American companies with at least \$1 billion in revenue ^[4]; its Finance Trends survey covers roughly 1,300 to 1,500 global finance leaders ^[8]. The FP&A-practitioner surveys — AFP's benchmarking survey ^[1], Vena's reports with BPM Partners ^[9] ^[10] — reach further down-market. Where a figure comes from an enterprise sample we say so, and our reading of what it implies for the mid-market is labelled as ours.

01 The state of practice: the spreadsheet is the system

Every prediction that Excel would leave finance has been wrong, and the

2025–2026 data is the most emphatic yet.

SPREADSHEET DEPENDENCY AND WHERE THE TIME GOES · PUBLISHED FIGURES, 2020–2026

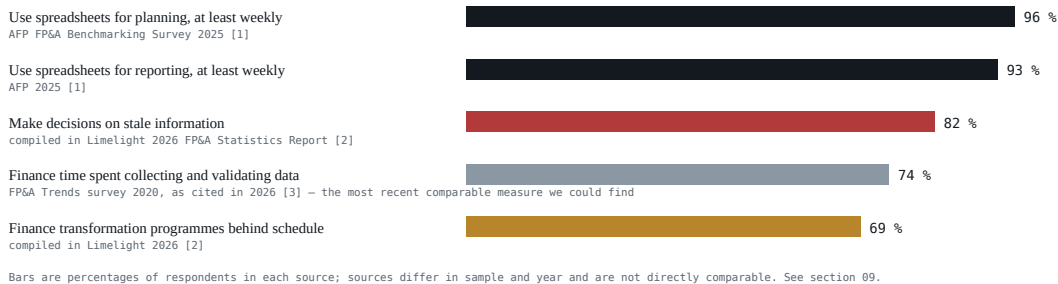


Figure 1. The practice baseline. Near-universal spreadsheet use is not a failure of adoption; it is the environment any planning platform has to fit into.

Two things follow from the AFP figures that vendors tend to skip. First, "adopting a planning platform" in practice means moving the *governed* part of the process — the model, the versions, the approvals, the numbers of record — off the spreadsheet, while planners continue to think, prototype and check in Excel. Platforms that fight this lose. Second, the spreadsheet's persistence is the reason the data chain is manual: if the plan lives in a workbook, the ledger reaches it by export. The 82% stale-decision figure ^[2] and the three-quarters-of-time-on-data figure ^[3] are the same finding measured twice.

The Vena 2026 FP&A Impact Report, from a survey of more than 400 finance leaders, describes teams generating valuable insight but constrained by data issues and manual processes, and businesses struggling to convert insight into action ^[10]. That is the practitioner's summary of the whole section: the analysis is often fine; the plumbing between the systems is not.

02 The forecasting gap

CFOs name forecast accuracy as their first cost-control problem, and the move to continuous forecasting is the response most of them are making.

61% of CFOs identify inaccurate forecasting as the single biggest barrier to controlling costs ^[2]. Industry commentary commonly reports forecast-to-actual variances of 30 to 40% for teams on manual processes ^[11]; we treat that as an indicative range from practitioners rather than a measured benchmark, but it matches what we have seen in mid-market implementations, where the first driver-based forecast typically halves the prior variance simply by tying volume to price and headcount to start dates.

The structural response, visible across the 2026 trend reporting, is the rolling forecast: a 12-to-18-month horizon updated monthly or quarterly using driver-based models, with scenario analysis as a standing feature of the monthly review rather than a crisis exercise ^[12] ^[13]. Our companion article on rolling forecasts sets out the mechanics; the landscape point is that the practice has moved from progressive to standard in the published guidance, while the tooling most mid-size companies run — a workbook per department — cannot support it past two cycles.

What accuracy problems usually are

Not model choice. In our experience the three causes, in order: the plan was built on a stale ledger; drivers were replaced by percentages; timing (start dates, collection days) was averaged away. Each is a data-chain problem before it is a forecasting problem.

What the surveys agree on

Forecast accuracy, data quality and manual process appear together in every 2026 source we read. Deloitte's enterprise CFOs rank data quality, access and usability (52%) alongside AI agents as the top transformation priorities ^[4].

What we would measure

Version-against-actual by driver and owner every quarter, and machine-versus-planner accuracy where a statistical baseline exists. Accuracy improves when it is attributed; the surveys show it is rarely measured at all.

03 The close and the data chain

The close is where the planning data chain starts, and the gap between

average and leading is wide.

MONTH-END CLOSE • DAYS • 2026 BENCHMARK AS REPORTED



The Hackett Group 2026 benchmark, as reported by a secondary source [14]. We could not consult the primary study; the ratio is consistent with what we see in mid-market clients that have automated the ledger-to-plan chain, and inconsistent with nothing else we found.

Every day between close and load is a day the plan is measured against numbers that have already moved.

Figure 2. The close. A six-day average close followed by a manual export means the plan sees actuals around working day eight to ten; a two-day close followed by a nightly load means day three.

Our interest in the close is downstream. The planning data chain — ledger to load to model to statement — begins when the books are closed, and each manual link adds days. Half of large-company CFOs name digital transformation of finance as their top 2026 priority and 49% make automating routine work their top talent priority ^[4]; the mid-market equivalent is smaller in scope and larger in relative effect, because the chain is shorter and the person who runs it is usually the same person who is supposed to be analysing the result.

04 AI: intent, adoption, use

The most consistent pattern in the 2026 data is a ladder with a wide gap between its top and bottom rungs. Nearly every CFO believes AI matters; a

minority run it in a core workflow; fewer still in forecasting.

THE AI LADDER · INTENT AT THE TOP, USE AT THE BOTTOM · 2025–2026 SOURCES




Figure 3. The AI ladder. The top two rungs are enterprise CFO sentiment; the middle is broad adoption "in some form"; the bottom is what changes a number in a plan. The distance between the top and the bottom is the finding.

What sits between intent and use

Three things, on the evidence. *Data.* Deloitte's enterprise CFOs put data quality, access and usability (52%) level with AI agents as a transformation priority ^[4]; a model cannot forecast from a ledger it cannot reach. *Security and governance.* In Deloitte's Finance Trends survey of 1,326 global finance leaders, data-security concern is the number-one barrier to AI adoption at 47%, and in a third of organisations the chief information security officer owns AI governance ^[8]. *Starting point.* 68% of CFOs say they have been slow to adopt because they do not know where to start, and the most-used tool is a general-purpose assistant rather than finance-specific software ^[5]. L.E.K. finds about a quarter of finance functions using AI features inside third-party software, with 44% planning to within three to five years ^[15].

The practitioner's reading

The ladder is not a maturity curve that time will climb on its own. The rungs are separated by architecture: an AI feature that reads a governed model, as a named user, and files a finding a person dispositions, is a different object from a chat window with a pasted spreadsheet. Companies stuck at "some form" have usually adopted the second and are waiting for it to become the first. Our companion articles on MCP and governed agents set out what the first looks like; the landscape point is that the 17% got there by design, not by experimentation.

05 The tooling market: price, time, and the three tiers

The planning-software market a mid-size company shops in has three tiers, and the published evidence on what each costs and takes is more consistent than the vendors' pricing pages.




Figure 4. Price against time. The mid-market tier is affordable and reasonably quick, and buys planning alone; the integration and analytics layers are separate purchases, which is where the total cost and the implementation calendar actually go.

TIER	EXAMPLES, AS COMMONLY SHORTLISTED	WHAT IT DOES WELL	WHAT THE BUYER STILL HAS TO ADD
Excel-centric FP&A	Datarails, Vena, Cube	Governance and workflow on top of workbooks planners already use; low friction to adopt	Integration beyond one or two ERPs; multi-entity, multi-calendar, currency; optimisation and forecasting engines; an analytics layer
Mid-market planning	Planful, Workday Adaptive Planning, Pigment	Modelled planning with versions, workflow, consolidation in the upper tier; two-to-four-month deployments	An iPaaS for the third source onward; BI for the dashboard; add-on modules for advanced features; partner services
Enterprise	Anaplan, Oracle Cloud EPM, SAP, OneStream	Scale, consolidation depth, connected planning across functions, partner ecosystems	Six to twelve months and a partner; licence and services sized for companies ten times larger
Tier assignments follow published comparison guidance [7][16]; capability descriptions are practitioner summaries and should be verified against each vendor's current release.			

Two market facts frame the mid-market buyer's decision. First, the 2025–2026 budget data shows money moving toward planning technology: one compilation reports 65% of CFOs increasing FP&A technology budgets by at least 20% in the year, with 70% of teams on cloud platforms ^[13]. Second, 69% of finance transformation programmes are reported behind schedule ^[2]. Spend is rising faster than delivery — which, in our reading, is what happens when a company buys the mid-market tier and then discovers the two or three products it has to add to make it a system.

06 What the leaders do differently

The surveys describe their own leading cohorts. Read together, the descriptions converge on three practices and one attitude — none of which

depends on company size.

PRACTICE 1

One planning data core

Leading teams "centralised their planning data" and build "unified data cores" before AI ^[2]^[5]. In architecture terms: one copy of actuals, drivers and plan versions, in one store the models and the reports both read.

PRACTICE 2

The chain automated

They "automated the manual work" ^[2] — the export, the load, the reconciliation, the distribution — and 49% of enterprise CFOs make exactly this their top talent priority ^[4], because it is where analysts' hours go.

PRACTICE 3

AI on named problems

They "put AI on real forecasting problems rather than treating it as a science project" ^[2], and 48% of strategy-leading finance functions have deployed specific agents ^[8]. Variance, headcount versus plan, forecast baselines — not a chatbot with a general remit.

ATTITUDE

Continuous, not annual

Rolling horizons, scenario planning as a standing fixture, driver models updated monthly ^[12]^[13]. The budget becomes a target; the forecast becomes the plan.

The practitioner's observation is that these are ordered. The data core comes first because the chain cannot be automated into three stores; the automated chain comes second because AI applied to stale data produces confident nonsense; and the named AI problems come third because they are only tractable once a governed model exists to point them at. Companies that start at the third step — most, on the CFO Connect evidence ^[5] — end up in the 45% who are still piloting.

07 A readiness self-assessment

Ten questions a leadership team can answer in an hour. Practitioner-authored; the levels describe what we see in mid-market finance functions, not a

published maturity model.

QUESTION	LEVEL 1 · MANUAL	LEVEL 2 · ASSISTED	LEVEL 3 · CONNECTED
How do actuals reach the plan?	Export and paste, monthly	Scheduled import into a planning tool	Nightly governed load into the store the plan reads, with a log
Where does the plan of record live?	A workbook per department	A planning tool, with Excel alongside	Named, locked versions; Excel for thinking, the platform for the number
What happens after a reorganisation?	History restated by hand, or not	Hierarchy updated; history frozen	A dated hierarchy version; as-was and as-is reporting
How is revenue forecast?	Last year plus a percentage	Volume × price by segment	Driver-based with a statistical baseline and intervals
How is payroll planned?	Headcount × average salary	By role, with bands	By position, with start dates, ramp, load and merit dates; HRIS synced
How far ahead can leadership see?	To fiscal year-end, shrinking	Twelve months rolling	Eighteen months rolling, re-cut quarterly
Who may change a number?	Anyone with the file	Users of the tool	Rules by entity and cost centre; locks by version; every cell audited
How is a variance explained?	Someone investigates when asked	A variance report	A finding raised automatically, with severity, dispositioned by a person
How does AI touch the plan?	Pasting into a chat assistant	Vendor AI features in the tool	Governed readers: tools exposed opt-in, run as a user, audited

How long did the last change take?

A quarter

Weeks, with a partner

Days, from a spec, promoted from a sandbox

Most mid-size companies we meet score level 1 on the first three rows and level 2 on the rest. The order of the rows is the order in which to fix them.

08 What this means for a company of 100 to 1,500 people

Practitioner opinion, stated as such, and tied to the evidence above.

The binding constraint has moved

Ten years ago the mid-market could not afford connected planning; the enterprise tier was the only tier that had it. The 2026 evidence — mid-market contract medians around \$35,000–\$40,000^[6], cloud platforms at 70% of teams^[13] — says capability and price are no longer the constraint. Implementation time and governance are.

Buy the chain, not the tool

The 69% behind-schedule figure^[2] is, in our reading, mostly the cost of assembling a system from three products. A platform that holds integration, planning and analytics on one store removes the seams that make transformation programmes late.

Keep the spreadsheet; move the number of record

96% weekly spreadsheet use^[1] is not going to zero. The realistic target is that the governed number — the locked version, the approved headcount plan — lives in the platform and Excel remains the analyst's scratchpad and the manager's inbox format.

Do the three practices in order

Data core, automated chain, named AI problems. The first two are weeks of work for a mid-size company on a single-application scope; the third is where the 87% intent^[4] becomes the 17% use^[5] — and it should be governed from the first agent, because the security concern that leads Deloitte's barrier list^[8] is legitimate.

Measure forecast accuracy, or nothing improves

61% of CFOs name it as their first cost problem ^[2]; almost none of the surveys report anyone measuring it by driver and owner. Version-against-actual, attributed, every quarter, is the single practice with the best ratio of effort to effect we know.

Expect weeks, and hold the scope to earn it

Published mid-market deployments run two to four months ^[7]. Four weeks is achievable — the companion playbook sets out how — but only for one application, one ledger, three sources and one approval chain. Multi-entity elimination and currency are a second wave, everywhere, on every platform.

09 Method and sources

No primary survey was conducted for this report. Figures are taken from the sources below as published or as reported by the secondary source named;

sample sizes and populations are given where the source publishes them.

REF	SOURCE	POPULATION AND DATE	USED FOR
[1]	AFP FP&A Benchmarking Survey 2025 (as cited by Datarails, Jul 2026, and Tesorio, 2026)	FP&A professionals; 2025	96% spreadsheets for planning weekly; 93% reporting; 23% AI in forecasting regularly
[2]	Limelight Software, 2026 FP&A Statistics Report (press release, 26 May 2026); a compilation of 50+ third-party benchmarks including PwC and SAP Concur	Compiled; 2026	61% CFOs forecast accuracy; 82% stale decisions; 69% programmes behind; 28% + 39% AI in forecasting (PwC)
[3]	FP&A Trends survey 2020, as cited by Houseblend (May 2026)	Finance professionals; 2020	74% of time on data collection, generation and validation
[4]	Deloitte, Q4 2025 CFO Signals (press release 13 Jan 2026; CFO Dive, 14 Apr 2026)	200 CFOs, North America, ≥ \$1B revenue; Nov–Dec 2025	50% digital transformation top priority; 49% automation top talent priority; 87% AI important; 54% AI agents; 52% data quality
[5]	CFO Connect, State of AI in Finance 2026 (Mar 2026)	Finance teams, Europe-weighted; 2026	56% adoption; 45% limited pilots; 17% core workflows; 68% do not know where to start; general-purpose assistant most used
[6]	Vendr deal data as reported by Aleph, "How much does FP&A software cost?" (Aug 2026)	Anonymised contracts; Aug 2026	Median annual contracts: Datarails \$33.3k, Vena \$34.3k, Planful \$42.4k, Pigment \$74.0k, Anaplan \$114.6k; implementation \$0–\$40k+

[7]	ChatFin, "AI FP&A Software 2026" comparison (Jun 2026)	Vendor comparison guidance	Deployment ranges by tier; enterprise all-in \$200k–\$500k+
[8]	Deloitte, Finance Trends 2026 and Q2 2026 CFO Signals (Jul 2026)	1,326 global finance leaders; 200 CFOs	47% data security as top AI barrier; 33% CISO owns AI governance; 48% of strategy-leading functions deployed agents
[9]	Vena / BPM Partners, 2025 State of Strategic Finance (Business Wire, 29 Apr 2025)	200+ senior finance professionals, North America; 2025	57% of finance teams using AI; data silos and manual workflow as constraints
[10]	Vena, 2026 FP&A Impact Report (Jun 2026)	400+ finance leaders, global; 2026	Insight-to-action gap; data and manual-process constraints
[11]	Tesorio, cash-flow forecasting comparison (2026)	Industry commentary	30–40% forecast-to-actual variance for manual teams — treated as indicative
[12]	FP&A Consulting Inc., "Future of FP&A: 2026 and beyond" (Dec 2025)	Practitioner guidance	12–18-month rolling forecasts as the 2026 standard; scenario planning as default
[13]	Limelight, "7 FP&A Trends for 2026" (Jun 2026)	Compiled	65% of CFOs raised FP&A tech budgets $\geq$ 20%; 70% on cloud platforms; continuous forecasting trend
[14]	The Hackett Group 2026 close benchmark, as reported by ChatFin (Jun 2026)	Secondary; primary not consulted	1.8-day vs 6.2-day close
[15]	L.E.K. Consulting, as reported by accoTOOL (Jun	Secondary	~25% use AI features in third-party software; ~44%

	2026)		plan to within 3–5 years
[16]	Aleph, "Best FP&A software, Q3 2026"	Vendor comparison	Tier positioning of Planful, Vena, Workday Adaptive
Where a figure is cited through a secondary source, the secondary source is named and the figure should be verified against the primary before being re-used. Figures from enterprise samples ([4], [8]) are applied to the mid-market as opinion, not as measurement.			

What would improve this report

A primary survey of mid-market finance teams — 100 to 1,500 employees, by name — on the ten readiness questions in section 07. None of the sources above segments by that band. We intend to run one for the 2027 edition and will publish the sample and the instrument.

EAConnect Planning · Research library · Hours, Not Months · Flagship

A synthesis of published surveys and market data, 2025–2026, with practitioner interpretation labelled as such. No primary research. Competitor prices are third-party medians, quote-based and subject to change.

3,579 words · 16 min read · — pages