BUYER'S GUIDE

The agentic *readiness gap*

A buyer's framework for judging any FP&A vendor's AI claims against what agent-driven finance requires.



FOREWORD

Foreword



By Christina Ross
Founder & CEO, Cube

Every FP&A vendor now has an AI story. Most of the demos look good. A clean question goes in, a confident answer comes out, and a chart appears.

As a three-time CFO, I learned to distrust any answer I couldn't trace. The question I'd ask about an AI demo is the one I asked about every number that crossed my desk: where did this come from, and would I put my name on it in front of the board?

Agents raise the stakes on that question. A single AI feature that drafts commentary can be checked by a person. A team of agents that monitors the business, drafts entries, builds scenarios, and assembles the board deck produces far more output than any person can recheck line by line. Whether that output can be trusted depends on what sits underneath the agents, and most of that is invisible in a demo.

This guide gives finance leaders a way to look underneath. It names six requirements for agent-driven finance, describes where most platforms stand on each today, and lists the questions that separate a real capability from a good demo. Use it on us and on everyone else you evaluate.

"AI without governance is a faster way to get the wrong answer."

Christina

Contents

CHAPTER 1	Why agents change the buying decision	4
CHAPTER 2	Requirement one: governed financial context	5
CHAPTER 3	Requirement two: trace to the source	6
CHAPTER 4	Requirement three: finance owns the rules	7
CHAPTER 5	Requirement four: a continuously current model	8
CHAPTER 6	Requirement five: reach into every surface	9
CHAPTER 7	Requirement six: people stay in control	10
SCORECARD	The readiness scorecard	11
CLOSING	About Cube The agentic FP&A platform and how finance teams use it	12

CHAPTER 1

Why agents change the buying decision



For most of the last decade, FP&A software was evaluated on modeling flexibility, reporting, integrations, and ease of use. Those still matter. Agents add a new test, because an agent is only as reliable as the financial foundation it works from.

A general-purpose AI tool pointed at raw data will answer a question about gross margin by reconstructing gross margin from whatever tables it can find. It might pick the right accounts. It might include a reclass that finance reversed last month. It will sound equally confident either way. AI can be 80% right. In finance, that's 100% wrong.

That's the readiness gap: the distance between the AI a platform can demo and the foundation agent-driven finance needs underneath it. The gap shows up in four places across the market today.

PLATFORM TYPE	WHAT IT DOES WELL	WHERE IT STOPS FOR AGENTS
Spreadsheet-bound	Flexible and familiar	The logic lives in files that agents can't govern or trust
Browser-bound	Modern and easy to adopt	Insight stays inside an app business leaders rarely open
Legacy platforms	Governed and scalable	Built around the annual plan, with changes that wait on administrators
General-purpose AI	Fast, capable answers	No governed financial context to trust the answers

Each chapter that follows covers one requirement, what most platforms offer today, what agent-driven finance needs, and the questions to ask.

CHAPTER 2

Requirement one: governed financial context



What agent-driven finance needs. Agents need to know what the numbers mean: how accounts roll up, which definition of ARR the company uses, how intercompany eliminations work, what drives headcount cost, and which adjustments finance has made. Data tells AI what the numbers are. Financial context tells AI what the numbers mean. That context has to be shared, so every agent and every person works from the same definitions.

Where most platforms stand today. Many AI features sit on top of a reporting layer or a raw data connection. They can read numbers. They don't carry the business logic, so each agent effectively rebuilds the model in the moment, and two agents can return two answers to the same question.

Questions to ask

1. **1.** When your AI answers a question about margin, which definition of margin does it use, and who set it?
2. **2.** If finance changes how a metric is calculated, does every agent pick up the change automatically?
3. **3.** Show me two different agents answering the same question. Do they match?

RED FLAG

The vendor's AI works from a data warehouse or export with no connection to the financial model finance maintains.

CHAPTER 3

Requirement two: trace to the source



What agent-driven finance needs. Every number an agent produces should open to the transactions behind it: which source system, which entries, which calculation, and which agent touched it. Auditors will ask. Boards will ask. Controllers need it to approve an agent's draft without redoing the work.

Where most platforms stand today. Many AI answers arrive as text with a summary of how they were produced, or with a link to a report. Getting from the answer to the specific GL transaction takes a person and several clicks, if it's possible at all.

Questions to ask

1. **1.** Pick a number your AI just produced. Take me to the transactions behind it.
2. **2.** When an agent drafts a journal entry or variance comment, what supporting evidence travels with it?
3. **3.** Can I see which agent produced an output and what data it used?

RED FLAG

The answer to "where did this number come from?" is a description of the model's reasoning instead of the transactions.

"They operate on patterns. You operate on facts."

CHAPTER 4

Requirement three: finance owns the rules



What agent-driven finance needs. As agents take on more of the work, control matters more. Finance should own the model, the planning process, the permissions, and the rules agents operate within: what an agent can draft, what needs approval, who signs. Finance should be able to change those rules without filing a ticket.

Where most platforms stand today. On many platforms, changing the model or the workflow requires an administrator, a consultant, or a services engagement. That was slow before agents. With agents, it means the rules lag behind the business, and finance loses control of the process at the moment it matters most.

Questions to ask

1. **1.** Can my team change a model, a definition, or an approval rule without your services team or ours?
2. **2.** Do agents inherit the same permissions as the person using them, or do they see everything?
3. **3.** Who configures what an agent is allowed to do, and how long does a change take?

⚠ RED FLAG

A meaningful model change requires a statement of work.

> **"The moment you need a consultant to change your model, finance lost control of the process."**

CHAPTER 5

Requirement four: a continuously current model



What agent-driven finance needs. Agents that monitor the business need a model that reflects the business now. If actuals refresh monthly and the forecast updates quarterly, an agent watching for variances is watching last month.

Where most platforms stand today. Most planning tools were built around a calendar: load actuals at close, reforecast on a schedule. Data refresh is often a batch job finance triggers by hand, and the forecast is a snapshot.

Questions to ask

1. **1.** How current is the data your agents see right now, and what determines the refresh?
2. **2.** When something changes in a source system mid-month, how long until the forecast reflects it?
3. **3.** Can your agents flag a variance before the close, or only after?

RED FLAG

Agent monitoring that only runs after month-end close.

CHAPTER 6

Requirement five: reach into every surface



What agent-driven finance needs. Decisions happen in pipeline reviews, Slack threads, board decks, spreadsheets, and AI assistants. An agent's work is only useful if it reaches those places, with the same governed context intact. That includes the external AI tools leaders already use, connected through open standards such as MCP.

Where most platforms stand today. Most AI features live inside the vendor's own application. To use the answer, someone logs in, finds it, and copies it somewhere else. Business leaders rarely log in, so the insight rarely reaches the decision.

Questions to ask

1. **1.** Can my team ask your platform a question from Excel, Google Sheets, Slack, Teams, and an AI assistant like Claude, and get the same answer each time?
2. **2.** When the answer leaves your application, do permissions and trace to the source travel with it?
3. **3.** Do you offer an MCP server, and what can it do beyond reading data?

RED FLAG

AI that only works inside the vendor's interface.

"Any surface. One financial truth."

CHAPTER 7

Requirement six: people stay in control



What agent-driven finance needs. Agents should monitor, analyze, model, investigate, and coordinate. People should judge, prioritize, decide, partner, and lead. A platform built for agent-driven finance makes that split explicit: drafts are labeled as drafts, approvals are enforced, and nothing an agent produces reaches the business without a named person signing off when the rules require it.

Where most platforms stand today. Vendor claims vary widely. Some promise fully autonomous finance. Others offer AI that can't do anything beyond answering questions. Neither serves a CFO who has to answer for every number.

Questions to ask

1. **1.** Which of your agent's actions require human approval, and who sets that?
2. **2.** What happens when an agent is wrong? How would I find out, and how fast?
3. **3.** Which agent capabilities ship today, and which are on the roadmap?

RED FLAG

"Autonomous" used without a clear answer to who signs.

cube

Cube was built around these requirements. Cube Operating Intelligence™ (CubeOI™) gives every agent the same governed financial context, with trace to the source on every number. Finance owns the model through a no-code experience, and agents work within the rules finance sets. The same context reaches Excel, Google Sheets, PowerPoint, Google Slides, Slack, Teams, the web, and AI assistants through the Cube MCP Server. We'll take the scorecard test with your data.

SCORECARD

The readiness scorecard



Score each vendor on each requirement: 0 if the capability isn't there, 1 if it exists with significant limits, 2 if the vendor demonstrated it live with your data or a realistic sample.

REQUIREMENT	TEST	SCORE (0–2)
Governed financial context	Two agents give the same answer to the same question	
Trace to the source	Any AI number opens to its GL transactions	
Finance owns the rules	Finance changes a model and an approval rule without services	
Continuously current model	Agents flag a variance mid-month	
Reach into every surface	Same answer in a spreadsheet, Slack, and an AI assistant	
People stay in control	Approvals are enforced and drafts are labeled	
Total		/ 12

How to read the score

10–12

Ready to put agents on real finance work.

6–9

Useful AI features. Expect to build controls and workarounds yourself.

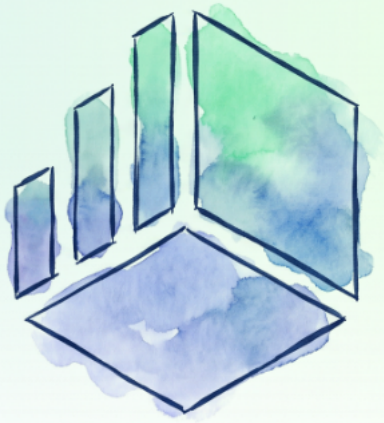
0–5

A demo layer. Agents will create more review work than they save.

Ask every vendor to score itself first, then run the tests. The difference between the two scores is worth discussing.

Ready to put this guide into practice?

Learn more about Cube



The agentic FP&A platform that helps you tell the story behind the numbers. Founded by three-time CFO Christina Ross.

Start your free trial →

Or get a demo

We believe the future of business is finance-driven. Cube connects your business data into one governed model and brings it everywhere decisions happen.

How finance teams use Cube

Close and consolidation

Multi-entity rollups and a faster month-end.

Budgeting and planning

Budgets, headcount, and revenue plans.

Forecasting and scenarios

Rolling forecasts and what-ifs in minutes.

Reporting and board decks

Packs in Excel, Sheets, and slides.

Variance analysis

Drivers explained and traced to the GL.

KPIs and performance

One set of metrics for the business.

AI Agents: Your Always-On Digital Finance Team

Charlie, the superagent, routes work to four FP&Agent teams. People stay in control.

The Data Managers
Integrate and reconcile

The Analysts
Analyze and report

The Planners
Forecast and model

The Business Partners
Shape decisions