



Executive
Perspectives

Applying Agentic AI in the Finance function for transformative impact

Finance

June 2026

Introduction

Agentic AI is the topic at the top of CFOs' minds today. Building on the previous edition of this series on AI in finance, we share our most recent insights from working with finance functions across industries. With AI agents now able to reason, plan, and execute multistep finance workflows with limited human input, the question for CFOs in 2026 is no longer whether AI will reshape finance but how fast - and how broadly to act.

The signal from CFOs is unambiguous, yet most finance functions remain at the copilot stage and point automation, applying AI with limited reimagination of the underlying processes, roles, and controls.

In this edition, we discuss how agentic AI is reshaping the finance function and how CFOs can capture the full value. This edition provides a guide for CFOs and finance leaders on how to:

- 1 Recognize the shift to agentic AI and where it creates new value in finance
- 2 Reimagine end-to-end workflows so agents execute and humans own judgment and exceptions
- 3 Build the five enablers - process, tech, data, operating model, and governance - that make agentic finance scale
- 4 Sequence the journey from focused pilots to an always-on, autonomous finance function

**In this BCG
Executive Perspective,
we outline our vision of
the future of finance
with agentic AI**



Executive summary | Agentic AI is rapidly evolving, and demonstrating tangible value with the right processes and people enablers

Agentic AI is rapidly evolving, driving urgency for CFOs

- **AI has moved from a productivity tool to a strategic imperative:** 88% of CFOs view it as essential, and 96% expect transformative impact
- **The 2030 vision is always-on finance:** real-time close, continuous forecasting, and autonomous optimization of capital and business performance
- New agentic capabilities can fundamentally shift finance's role—**from a producer of numbers to an architect of value**—enabling the function to perform faster, better, and more efficiently

Pioneers are already realizing significant value

- The gap between leaders and followers is **widening fast; pioneers are running multistep agentic workflows—not just pilots—across FP&A, accounting, and treasury**
- **The results are structural, not incremental:** AI pioneers are unlocking 90% reporting automation, more than 80% touchless invoice entry, and more than 30% capacity freed for higher-value work

Successful AI adoption requires reimagination of processes, tech, and governance

- **70% of the transformation is from processes and people reimagination:** point fixes don't compound; winners reimagine workflows, roles, and controls across the full value chain
- **The tech has fundamentally shifted:** agentic coding, AI coworkers, and autonomous workflows make what was theoretical now deployable at greater speed than before
- Model **guardrails and human oversight** are critical to ensure the appropriate and reliable use of AI and to ensure confidence in deploying AI solutions for finance



Section 01

Context and evolution of AI in finance

The CFO agenda has shifted—AI is now a strategic imperative, not an experiment



Prioritizing AI

~88%

CFOs who see AI as essential or important in their finance priorities

Expecting impact

~96%

CFOs who expect a significant or transformative impact from AI in the next five years

Realizing ROI

>25%

ROI realized by leading finance organizations that have implemented AI at scale

The art of the possible has evolved with the rise of agentic capabilities—finance can now move from modeling to autonomous execution

Currently available and deployed capabilities

Evolving capabilities

Powered by agentic AI

	Modeling and analysis	Reporting and insights	Recommendations	Orchestration	Autonomous execution
Description	Retrieve, aggregate, and model financial data; predict trends; analyze external variables	Dynamically generate reports, charts, and second-order insights to enable decision making	Generate deep insights (down to transaction data), with suggested actions and financial impact	Orchestrate multistep workflows across systems with human oversight and rule-based and judgment logic	Observe, plan, assess, and execute E2E decisions autonomously within guardrails
Example use cases	Automated refresh of budgets and forecasts; scenario modeling	AI-generated variance commentary and board reporting	Journal entry anomaly detection and resolution; performance insights	Real-time cash and FX optimization agent; procure-to-pay processing agent	Dynamic discounting; proactive intervention for timely collections
Tech	Machine learning first-generation LLMs	GenAI	Coding agents and copilots	Multiagent workflows	Agentic automation, including judgments

Tech solutions compound and are **not mutually exclusive**; optimal solutions require tailoring and **will require layering of different technologies** (e.g., ML models and AI agents)

Finance AI as unique design requirements, and getting them right separates leaders from laggards

Finance complexities

Controls and security

- High bar set by controls and regulatory standards (e.g., SOX, GAAP, and IFRS)
- Segregation of duties controls must be preserved in agentic design

Accuracy and audit readiness

- Understanding of calculation and drivers essential for decision making
- Need for audit trail and transparency across workflows

Data access and consistency

- Finance accountable for single source of truth and quality of financial data
- Finance owns calculation logic for statutory, management, and investor financial statements

Change management

- Finance teams usually conservative and not necessarily inclined to trust AI outputs
- Finance teams often slow to experiment with GenAI and no-code tools for their daily work

AI design considerations

- Monitor model integrity and compliance with automated error and risk thresholds
- Maintain control by defining agent objectives and building human-in-the-loop workflows

- Ensure model explainability (avoid black box), with ability for human to override
- Rules and triggers (e.g., policy documents) for workflows must be clearly referenced

- Set guardrails to restrict AI from overwriting source data or critical logic
- Monitor for shadow capabilities building alternate or nonstandard outputs

- Encourage citizen-developer experimentation by providing tools, launching certifications, and creating superuser and black belt programs
- Create cross-functional teams for big builds

Note: SOX = Sarbanes-Oxley; GAAP = generally accepted accounting principles; IFRS = international financial reporting standards.

Finance's North Star is to evolve from producing numbers to architecting enterprise value

AI unlocks a new art of the possible across finance...

Planning, budgeting, and forecasting

Ecosystem of AI agents builds and updates dynamic, driver-based plans, continuously refreshing forecasts and recommending corrective actions in real time

Reporting and business intelligence

Always-on intelligence engine delivers real-time financial and strategic insights, with proactive recommendations

General accounting

Agentic close autonomously collects, matches, and reconciles data, flagging exceptions and driving continuous real-time close

Finance operations (e.g., S2P and O2C)

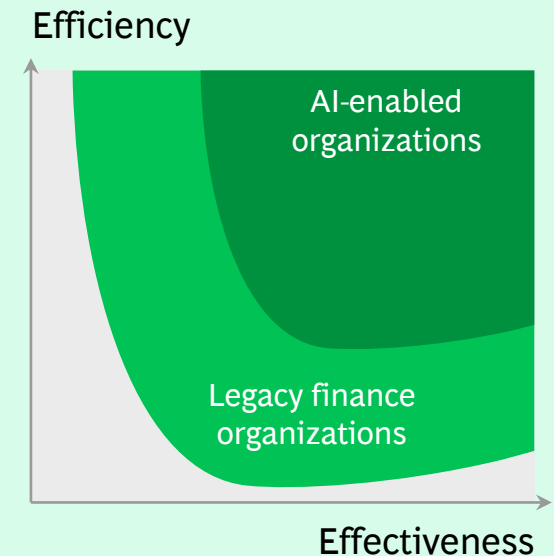
AI agents autonomously orchestrate E2E transactional workflows, such as gathering and entering data, running checks vs. policies, routing approvals, resolving exceptions, and communicating with suppliers and customers

Expert functions (e.g., treasury, tax, internal audit, and IR)

AI agents continuously optimize liquidity, hedging, and working capital, proactively surfacing opportunities to reduce risk and unlock capital

Note: S2P = source to pay; O2C = order to cash; IR = investor relations.

...improving both efficiency and effectiveness...



... and elevating finance's role from a producer of numbers to an architect of value

Use cases | Key opportunities exist for AI across finance processes

Planning, budgeting, and forecasting	Business strategy: AI surfaces market signals and scenario implications	Driver-based plans/forecasts: AI agents continuously build and refresh dynamic forecasts	Variance investigation: GenAI performs root-cause analysis and builds commentary	Scenario modeling: AI recommends corrective actions and refines driver trees
	Reporting and business intelligence		Ad hoc reporting: Natural language querying, automated visualization, and on-demand insight generation	Proactive monitoring: Continuous alerts on overspending, delays, and working capital—with recommended actions
General accounting	Agentic close: Autonomously matches, reconciles, and flags exceptions and drives continuous real-time close	Consolidation and filings: AI-based reconciliation and statutory report preparation		Compliance and policies: AI controls, error detection, and GenAI policy writing
	Finance operations		Order to cash: AI agents orchestrating E2E, including prediction models for credit scoring, early warnings for DSO and aging, predictive collections, exception resolution, and customer communications	Payroll: Autonomous execution and controls
Expert functions	Treasury: Continuous liquidity, hedging, and working capital optimization; proactive risk and capital surfacing	Tax: AI provisioning calculation and proactive flags for deferred tax impacts	Investor relations: Earnings Q&A preparation and investor sentiment analysis	Risk management: Continuous pattern recognition, early warnings, and fraud detection
	Cross-finance opportunities: Conversational interfaces, guided agentic workflows, and human-in-the-loop decision support			

High
 Medium
 Low

Analysis of opportunities based on additional value that AI can unlock (beyond other digital tools).
 Note: DSO = days sales outstanding; Q&A = question and answer.

Early movers are already unlocking up to 30% value and could realize more than 50% by adopting agentic AI for the long term

Finance function	Anticipated capacity value unlock ¹	Drivers of value
Planning, budgeting, and forecasting	<p>25% +20%-30%</p>	> Sharper market foresight, richer scenario planning, and more confident capital allocation
Reporting and business intelligence	<p>~20%-30% +30%-40%</p>	> Faster insights enabling proactive top-line intervention and margin optimization
General accounting	<p>~10%-15% +15%-25%</p>	> Fewer errors and lower financial risk; continuous controls replacing periodic reviews
Finance operations	<p>~20%-30% +20%-30%</p>	> Better customer and vendor experience, faster collections, and enhanced spending control
Expert functions	<p>~5%-10% +5%-15%</p> <p> ■ Current and near-term value unlock ■ Long-term agentic uplift </p>	> AI-empowered processes; improved working capital position and unlocked cash, freeing capital for reinvestment
Long-term potential: 50%+ value unlock		

Note: Value unlock refers to efficiency gains from AI-enablement of processes; additional gains from interest income earnings, ROI from reinvestment, and so on.
 Source: 2025 BCG CFO & Finance Executive survey; BCG experience.



Section 02

Client impact stories

Case study 1 | AI and agents are deployed across finance teams to speed decisions and improve controls at a global technology company

From: Manual and siloed

Manual processes
consuming finance capacity
and limiting strategic output

Fragmented data and slow reporting cycles limiting decision speed

Long forecasting cycles
with limited scenario modeling capability

To: Streamlined, AI-first finance team

15%+
FTE capacity reduction across finance

Reduced manual work across all finance processes (e.g., from agents in S2P, FP&A, and treasury)

Further cost savings and avoidance using proactive insights from financial analysis, contracting, and expense-auditing agents

90%+
Reduction in reporting time

AI-automated financial reporting freed up time by extracting data and writing business partner reports and commentary

Driver and outlier analytics built instant insights for decisions

25%+
More accurate forecasting

ML-driven budget and forecast generation pre-seeded data and runs models, freeing capacity to analyze scenarios

Cash deployment agent optimized liquidity and executed actions to generate additional interest income

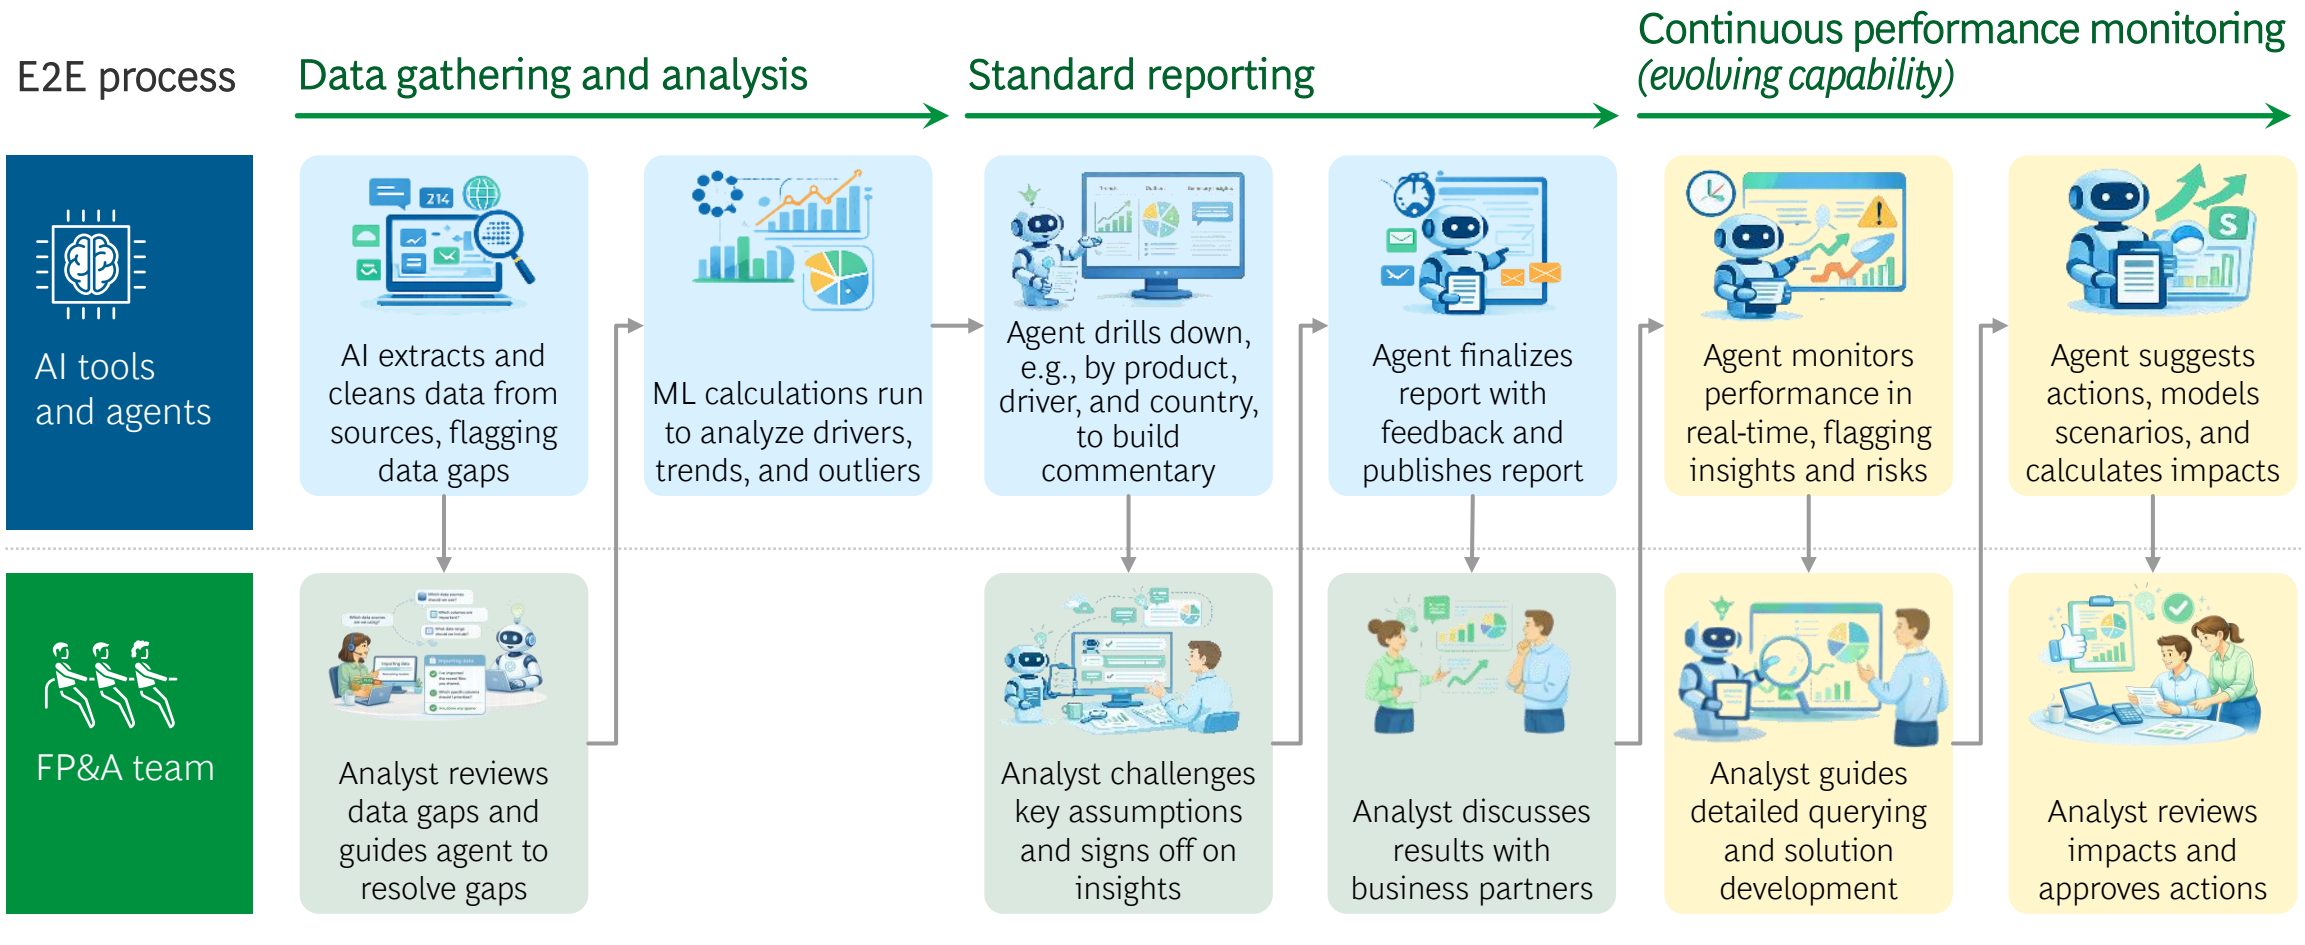
Use cases | The company implemented AI agents across the full spectrum of finance activities

Planning, budgeting, forecasting	Reporting and business intelligence	General accounting	Finance operations	Expert functions
Long-range plan build	QBR and recurring report generation	Transaction review and anomaly detection	AI-guided purchase operations (PR and PO)	Cash optimization and deployment agent
ML-driven plan and budget creation	Industry trends and signals monitoring	AI journal entry creation support	Touchless P2P invoice processing	Investor and market sentiment monitoring
ML-driven forecast generation	Market intelligence and peer benchmarking	AI-enabled reviews and reconciliations	Expense auditing	AI-driven FX hedging
AI-based variance analysis and reporting	AI-driven pricing optimization	Close orchestrating agent	Customer contract generation	Tax compliance agent
Agentic AI-enabled scenario modeling	Ad hoc analytics and investigation	Statutory report generation	Customer invoice generation	Tax optimization agent
	Use case deep-dive follows	Compliance and controls and policies monitoring	Predictive collections and automatic application	
			Vendor and customer help desk	

■ Prioritized use cases ■ Future use cases under consideration

Note: QBR = quarterly business review; PR = purchase requisition; PO = purchase order.

Deep dive | The reporting and business intelligence process was redesigned from manual data gathering to always-on insight



100% of standard reporting automated; 90% of analyst capacity redirected toward proactive business intervention

Analysts can now pursue advanced AI insights by training agent on business context and reasoning

Case study 2 | Implementing an FP&A insights agent enabled a large tech company to shift to real-time, on-demand business intelligence

From: Fragmented and slow

Inability to drill down from commercial outcomes to operational drivers

Slow to adapt to fast-moving business trends

Leadership decisions made without real-time data or forward-looking analysis

To: Agile and data driven

80%+
Faster ad hoc analysis

Increased ability to correlate data from driver tree models that create dynamic, live-edit links between key variables

Enabled deep-investigation analysis into source data, including extracting individual transactions, past trends, and outliers

2x-4x
Faster reporting generation

On-demand scenario modeling and planning enabled by BCG-developed, off-the-shelf Insights Agent tool

Enhanced data-driven decision making

Increased depth in analysis enabled by Insights Agent generating second- and third-order recommendations

Live charts that can be created during meetings and exported to decks enabled **data-informed leadership decision making**

Use case 1 | The agent transformed variance analysis from knowing what happened to automatically understanding why

The problem AI solved

Today, variance analysis is a manual, multiday exercise requiring FP&A analysts to stitch together data from multiple sources

How AI works

- 1 **Establishes baseline**
Agent calculates expected outcome vs. actuals for operational KPIs across the business (e.g., gross adds and churn)

- 2 **Identifies deviation**
Measures magnitude and direction of variance by driver and dimension (region, business unit, product, and customer cohort)

- 3 **Diagnoses root cause**
Traverses the driver tree to attribute variance to specific operational causes: pricing, volume, mix, competitive events, and external macro factors

- 4 **Recommends actions**
Generates prioritized, actionable recommendations—specific to market or business unit where driver is most addressable

Note: SMB = small and medium-sized business.

Example of AI-powered variance analysis

Analyst What's driving the revenue shortfall in North America this quarter vs. plan?

Agent Revenue is 8.2% below plan, driven primarily by volume (-6.1pp) in the SMB segment, with secondary impact from adverse product mix (-2.4pp). Pricing is neutral. Volume decline correlates with a competitor promotional campaign launched October 3rd in the Northeast region.

Analyst What actions should we take to course correct?

Agent Three options ranked by expected impact: (1) Targeted retention offer in Northeast—estimated +2.1pp recovery. (2) Accelerate Q4 pipeline in SMB segment—estimated +1.8pp. (3) Bundle promotion to shift mix—estimated +0.9pp. I can model the P&L impact of each scenario.

Response time: less than 30 seconds vs. two to three days for traditional analysis

Use case 2 | The company adopted AI-powered planning and forecasting, translating model outputs into real-time P&L impact for faster decisions

Use case examples

Rolling forecast refresh

Request: Update Q4 forecast with last week's actuals

Result: Agent refreshes all downstream nodes automatically

Competitive disruption scenario

Request: Model a competitor cutting price by 15% in our core market

Result: Generates P90, P50, and P10 revenue outcomes with confidence bands

Strategic initiative modeling

Request: Three-year EBITDA impact of expanding into two new regions

Result: Full P&L and cash flow simulation with sensitivity analysis

Macro stress test

Request: Simulate 200 bps rate increase and 2% GDP contraction on consumer revenue

Result: Multidriver propagation through full tree

Embedded AI modeling at the node level

Multivariate regression: Quantifies impact of multiple drivers on a KPI simultaneously

Time-series and predictive models: Captures seasonality, trend and calendar effects for rolling forecasts

Elasticity modeling: Assesses dynamic price-volume elasticity at product, segment, or market level

Stochastic and Monte Carlo methods: Determines probability distributions across key drivers for confidence intervals and risk-weighted scenarios

Driver correlation analysis: Uses ML-based detection of nonobvious relationships between operational metrics and financial outcomes

Case study 3 | A technology and SaaS company reimaged the procure-to-pay process to be an end-to-end touchless process with AI

From: Manual and risk averse

Error-prone processes due to legacy, inconsistent policies, manual reviews, and nonintegrated tools

Human-intensive workflows due to nonstandard processes, high exception volumes, and legacy tools

Highly intensive and slow vendor onboarding and management processes

To: Standardized and AI enabled

80%+
Touchless
invoice entry

Automated multistep tasks using bots and AI agents across systems (e.g., collecting, reading, invoice entry, and collation)

Eliminated PO checks and manual invoice processing steps by changing policies and implementing rules and ML decision logic

~45%
FTE
capacity unlock

Shifted human focus **from routine review to exception handling**, with AI agents running validations on standard process steps

Reduced nonstandard workflows (e.g., country-specific differences) and streamlined routine finance tasks

15+
Reduced days
for onboarding

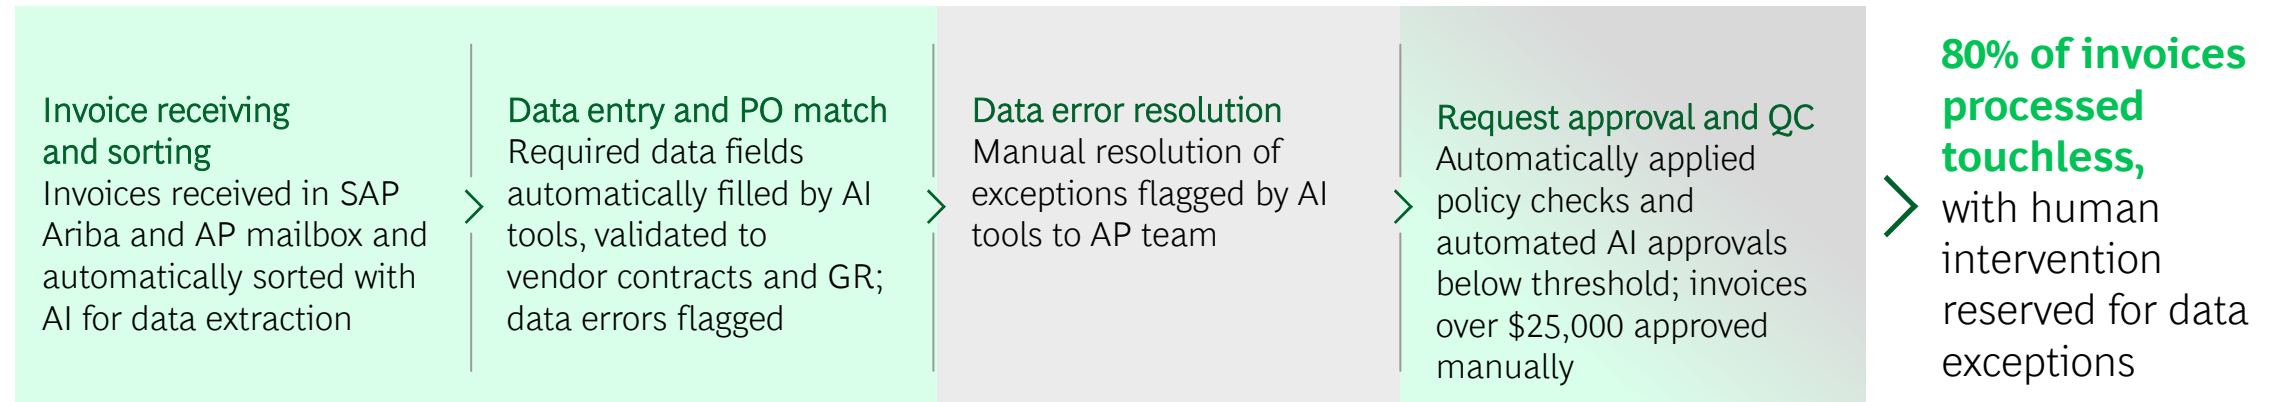
Redesigned internal policies to eliminate low value-added checks, raise thresholds, and restrict exception processing

Implemented **AI tools** on top of ERP and procurement systems for bank data validation and fraud checks

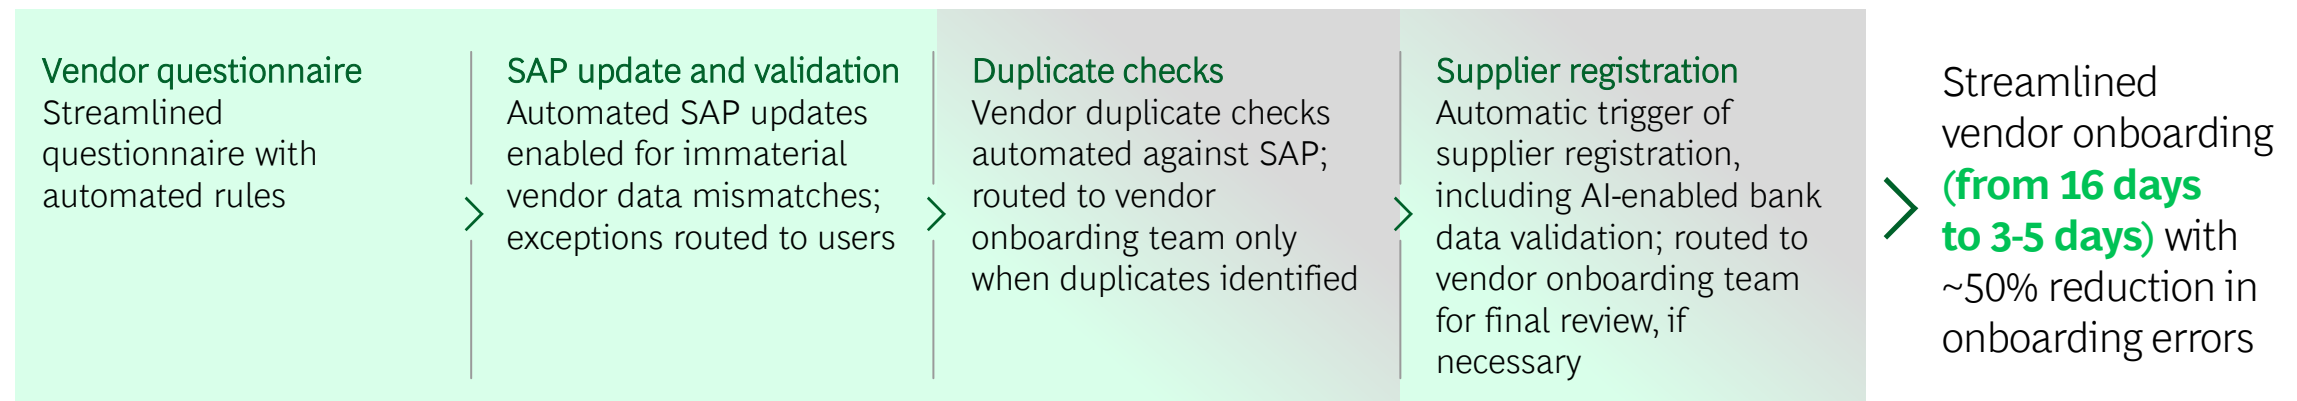
Note: ERP = enterprise resource planning.

Future process | The company is building agentic workflows to automate data entry, review, and approval, flagging areas for human action

Target state: Touchless invoice processing



Target state: Automated vendor onboarding



Manual Agent enabled with human in the loop

Note: AP = accounts payable; GR = goods receipt; QC = quality control.



Section 03

Getting started on an AI transformation journey

Winning with AI requires reinvention across five dimensions— simultaneously

1	Process reimagination	<ul style="list-style-type: none">• Drive top-down AI mandate, define bold finance vision, and build roadmap• Redesign E2E workflows with AI-first mindset
2	Tech ecosystem redesign	<ul style="list-style-type: none">• Deploy LLM and coding platforms to accelerate citizen development and prototyping—then layer orchestration platforms to shift to agentic workflows• Assess AI roadmap of existing tools in tech stack, assess build vs. buy options for prioritized workflows, and build and implement new tools where needed
3	Data foundations	<ul style="list-style-type: none">• Build the business context fabric: develop data relationships and driver trees, unify taxonomies, and clean and structure master data with GenAI• Transform data and build integrations and data structures for priority use cases
4	Operating model redesign	<ul style="list-style-type: none">• Redefine roles, OKRs, and incentives to align teams toward finance AI vision• Identify future skills and capabilities, and refresh talent and upskilling programs
5	Governance	<ul style="list-style-type: none">• Define guardrails for responsible AI (e.g., no-AI zones and human-in-the-loop reviews)• Define governing frameworks for AI roadmaps, prioritization, and investments

Note: OKR = objective and key result.

Process reimagination | Redesign end to end processes with AI, clarifying what agents execute and where humans intervene

Reimagined close and consolidation process with AI & Agentic workflows

Illustrative

E2E process

Always-on transaction processing and controls and reviews

Subledger and general ledger close

Consolidation and reporting



AI tools and agents



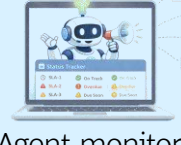
General accounting team



Agent runs journal entries overnight and creates report of escalated items



Agent monitors for policy and control breaches and anomalies




Agent monitors SL close deadlines, sends reminders, and escalates autonomously



Agent runs final adjustments (e.g., suspense account) based on rules



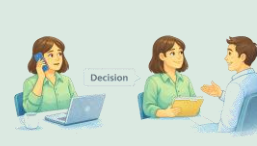
Agent consolidates and reconciles data, suggesting actions for mismatches




Agent drafts statutory reports, disclosure, commentary, and financial insights




Analyst reviews escalated and flagged items, applying business judgement to refine and approve accounting treatment




Analyst follows up or makes decision to postpone entry



Analyst reviews high-risk and high-value entries to approve GL close



Analyst assesses actions and signs off unreconciled balances



Analyst refines commentary with business and technical insights

Automation of journal entry preparation and reconciliation frees team capacity to proactively resolve anomalies and strengthen controls

Note: SL = subledger; GL = general ledger.

Process reimagination | Reimagining finance workflows with AI enables the Finance team to focus on value-adding work

Reimagined day in the life of an accounting analyst, under the new process

Illustrative

	8 a.m.	10 a.m.	12 p.m.	2 p.m.	4 p.m.	6 p.m.
Key tasks	<p>Review JE agent's overnight report</p> <ul style="list-style-type: none"> Review flagged anomalies and policy exceptions Apply business-specific accounting judgment and treatment 	<p>Clear anomalies or unprocessed transactions</p> <ul style="list-style-type: none"> Use Agentic tools to refine JE or calculation logic Validate proposed treatment against source data and transaction history 	<p>Validate journals and calculations</p> <ul style="list-style-type: none"> Review sensitive and/or high-value JEs Approve standard postings and escalate nonstandard items 	<p>Coordinate close and consolidation</p> <ul style="list-style-type: none"> Resolve reconciliation differences across entities and teams Review any escalations from autonomous close status tracking agent 	<p>Review trial balance commentary</p> <ul style="list-style-type: none"> Inspect flagged account movements, unusual variances, and eliminations Use Agentic tools to refine AI MDA commentary and reporting 	<p>Advise stakeholders and improve controls</p> <ul style="list-style-type: none"> Summarize key drivers and unresolved items for controller and FP&A Capture recurring root causes and flag control improvements
Lever	Monitoring and anomaly detection	Root-cause investigation and exception triage	Journal recommendation, escalation, and approval	Workflow orchestration and data aggregation	Variance review and commentary drafting	Follow-up action and continuous improvement
Savings	~15 minutes	~25 minutes	~25 minutes	~20 minutes	~20 minutes	~15 minutes

About 120 minutes of work saved per day can be reinvested in control enhancement, proactive issue resolution, and business partnering

Note: JE = journal entry; MD&A = management discussion and analysis.

Tech ecosystem redesign | Building a modular architecture enables shifting to an agentic finance function

Finance AI stack

Digital data platform (DDP) 3.0



1. Value-creation layer

Human experience channels
Ecosystem interfaces
Business services
Journey orchestration

^ Requests surface up

^ Actions pass down



2. Traditional and agentic AI layer

Orchestration services
Model and ecosystem garden
Agentic workflows
Guardrails management

^ Contextualized data

^ Queries and refinements

3. Data and knowledge layer

AI data services	
Semantic and knowledge	Data marketplace
Repository and storage	Data products
Distribution and integration	Data engineering
Data governance	

^ Governed records via APIs

^ Structured access

4. Core transaction layer

ERP	Systems of record
-----	-------------------

☆ Requires intentional redesign

Note: API = application programming interface.

Value-creation layer is human-centric by design

The full finance user journey—from spotting a variance to submitting a revised forecast—runs through natural-language prompts and smart dashboards, not raw data exports. Human review gates are embedded at every consequential step.

Traditional and agentic AI layer orchestrates intelligence E2E

Orchestrates both traditional AI and ML workflows and GenAI and agentic workflows. Controlled agent orchestration enables finance processes to run autonomously within governed guardrails, with human escalation paths built in.

Data and knowledge layer is a trusted, AI-ready data fabric

Provides the structured and semantic data foundation AI needs to reason about finance. Driver trees, business context, unified taxonomies, and clean master data are built here. Without this layer, AI outputs are unreliable and ungoverned.

Core transaction layer is an authoritative system of record

ERP and source systems remain the authoritative backbone. Clean APIs expose transactional data upward to the data and AI layers, enabling AI to act on real, governed numbers without bypassing controls.

Data foundations | AI doesn't just consume good data, it can also help to build it

Several manual tech and data barriers hinder finance today...

No clear data lineage, making E2E traceability and compliance reporting nearly impossible

No standardized source of truth, driving manual reconciliation and undermining reporting reliability

Fragmented data integration across finance and operational systems, creating persistent inconsistency

Lack of data ownership, governance, and access management, driving manual controls and review efforts

...but the barriers can be addressed with AI solutions...

GenAI-assisted lineage traceability: Define and automatically capture E2E lineage across pipelines; code parsing to pinpoint flow breaks

GenAI-assisted metadata enrichment: Assisted field definitions and standardization of naming conventions, schema, and taxonomy

Agent-assisted data integrity QC: AI agents monitor discrepancies, check data quality and gaps, and cross-reference to schema and standard definitions

Agentic-powered governance: AI agents pull and route data with defined workflows and automated review of data

...creating meaningful benefits for finance data and insights

➤ **Accelerated lineage compliance** with **enabled root-cause analytics** for exception management

➤ **Increased metadata accuracy** enables **reliable reporting and real-time** dashboards, reports, and business insights

➤ **Increased FTE capacity for insights and analysis** from reduced manual data management and reconciliation

➤ **Real-time data monitoring** dashboards with **high automation and protection** for sensitive data

Operating model redesign | Identify the roles that should be reinvented, and the skills that these roles require from Finance

Current finance team roles

Produce budgets and prepare forecasts >

Manually investigate variances and produce backward-looking performance reports >

Chase and reconcile fragmented data to investigate discrepancies and prepare adjustments >

Manually track cash positions and working capital drivers with limited predictive insight >

Perform broad, manual control testing and review high volumes of process exceptions >

Future finance team roles

Influence business leaders; dynamically shift capital, spending, and talent toward emerging growth signals detected by AI

Shape enterprise performance; translate AI insights into actions to improve financial and operational performance

Manage close exceptions proactively, with AI surfacing root causes and guiding controller and FP&A recommendations

Optimize liquidity with AI-driven foresight, enabling proactive funding, investment, and hedging decisions

Operate a live, self-learning control environment that continuously detects, prevents, and resolves risk

Illustrative

Skills the finance function needs to build



AI fluency to know how agents work, their limits, and when human oversight adds value



Deep business understanding to design optimization logic and performance interventions to guide autonomous agents



Strategic mindset to unlock new value pools, grow ROI from capital allocation, and enhance scenario and risk insights



Innovation and change mindset to continually reimagine roles and workflows as AI evolves, driving cross-functional improvement

Operating model redesign | AI doesn't replace finance functions, it collapses the handoffs, making teams faster, leaner, and more impactful



Today
Function based



The vision
Outcome based

Team shape

Layered hierarchies across corporate FP&A, accounting, and CoEs



Lean, flat pods: small teams orchestrate AI workflows and handling exceptions

Ways of working

Work handed off sequentially across FP&A, accounting, and CoEs, creating delays and misalignment



AI orchestrates across functions simultaneously: functions remain, sequential handoffs collapse

Time spent

Business unit-embedded generalists; most time spent on data gathering and reports



AI-augmented specialist squads; time freed for insight and business conversation

Scale

Scale people to scale output; more volume means more people and layers



Scale AI agents while keeping head count flat

Four actions to get to an outcome-based operating model

- 1 Restructure teams**
Flatten layers. Embed AI capability close to finance.
- 2 Redesign roles and responsibilities**
Humans own decisions; agents execute tasks. Make this explicit in every JD and process design.
- 3 Build new AI capabilities**
Upskill for orchestration and judgment. Build citizen developer communities and career paths.
- 4 Align OKRs and incentives to AI outcomes**
Rewire OKRs around outcomes and AI adoption, not process execution or report volumes.

Governance | Define governance principles that build trust and unlock the confidence to scale AI



Management accountability

AI executes and humans are accountable. Decision rights must be explicitly defined, with finance ownership of financial statements nonnegotiable.



Human oversight on high-judgment and high-risk items

Autonomous agents execute routine transactions and controls; humans approve material, complex, or unusual items to reduce risk.



Explainable and traceable outcomes

Any material agent output must be explainable, auditable, and supported by retained evidence.



Documented control design

Design, configuration, and monitoring must be strictly defined, documented, and tested; leverage rule engines to optimize performance within guardrails.



Data integrity, security, and privacy

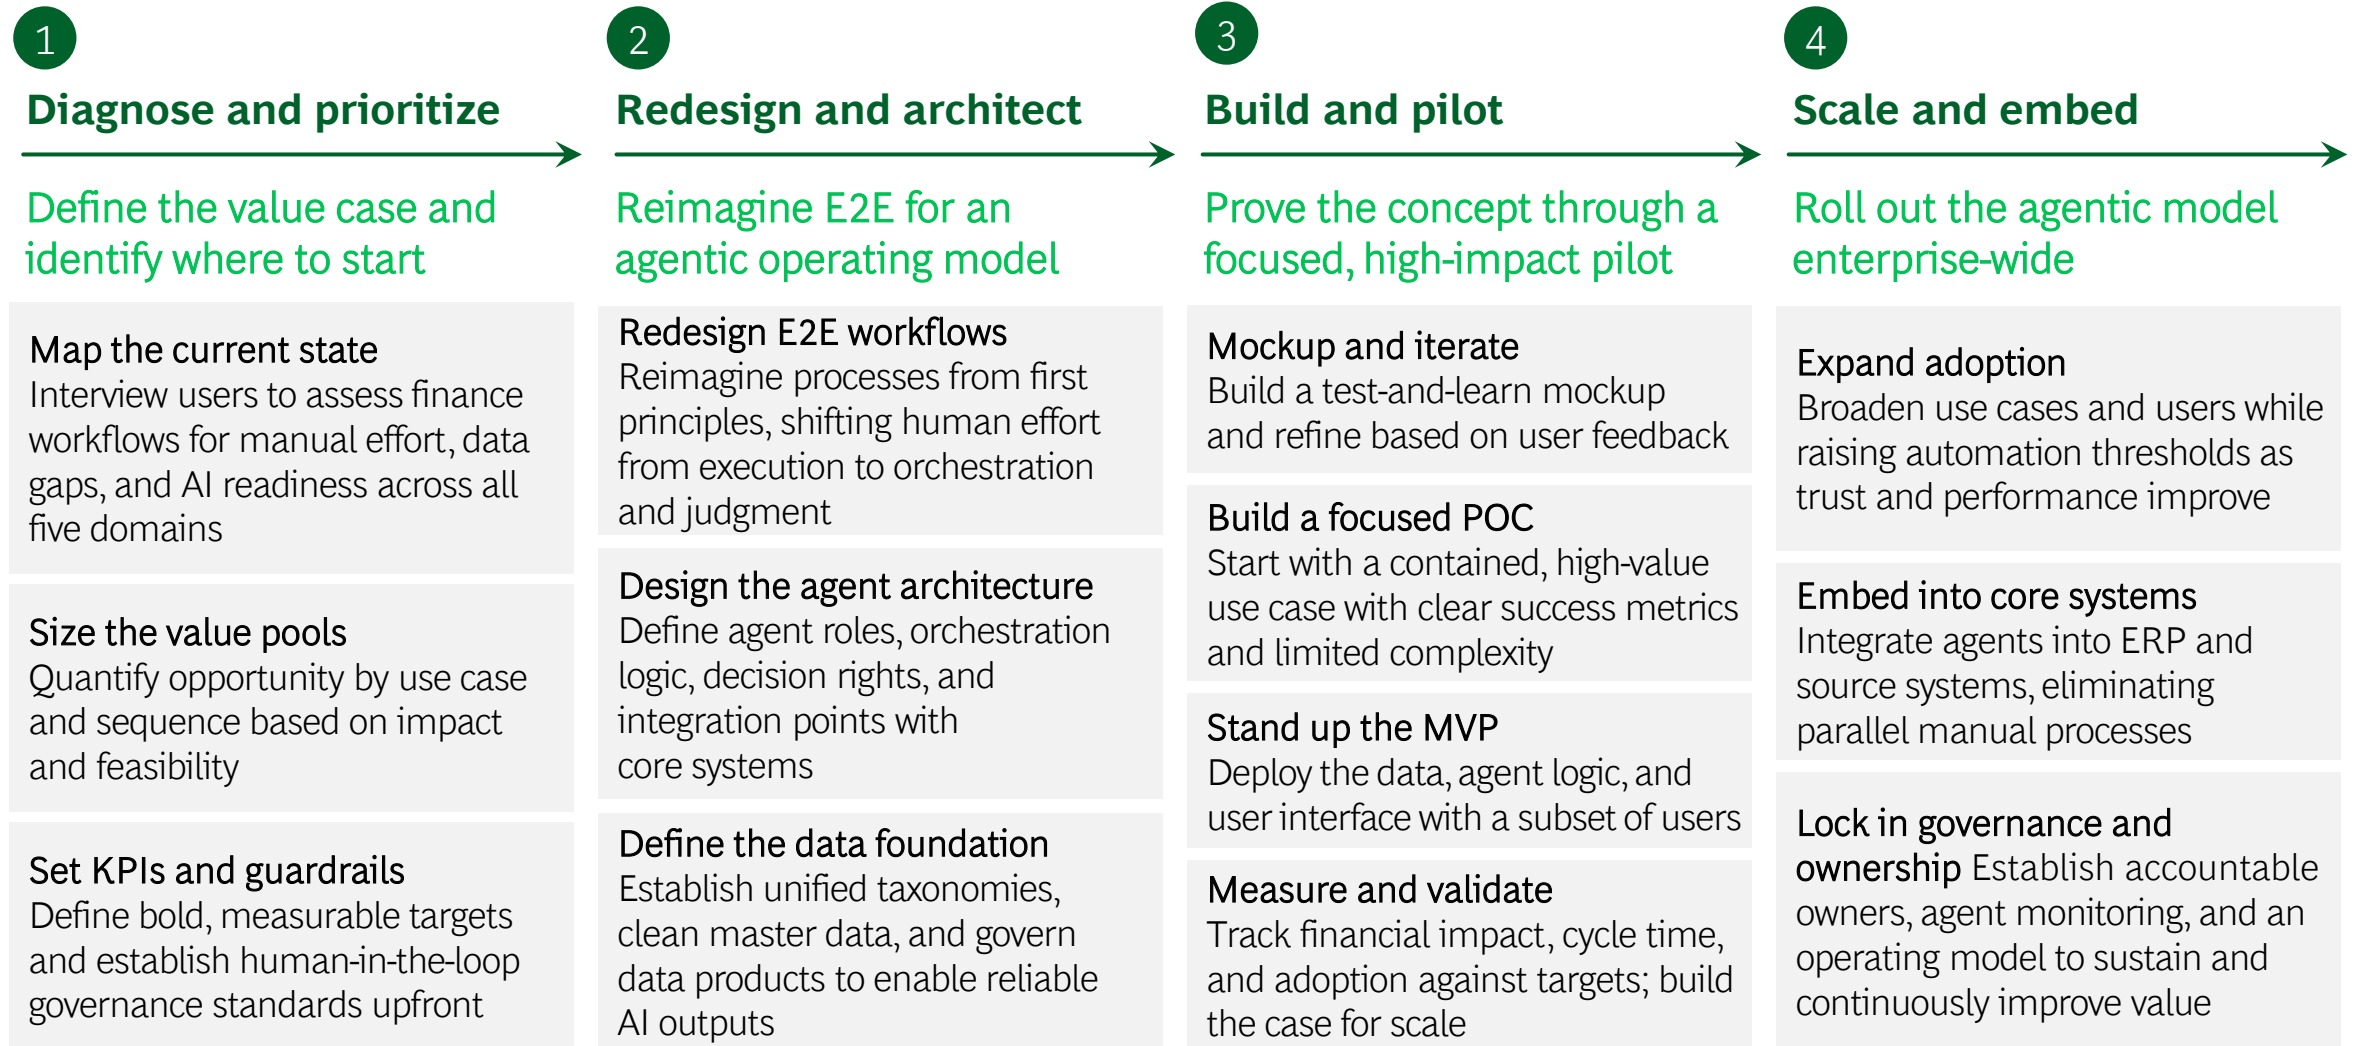
Lineage from source systems to agent inputs and outputs must be clear, with quality checks and access controls.



Model life cycle and change discipline


No silent updates; all changes go through testing and approval before production, with continuous monitoring of agent performance.

There are four key steps to redesign, pilot, and scale an AI-driven workflow



Note: POC = proof of concept; MVP = minimum viable product.

Algorithms and technology are only 30% of the equation. The real transformation is about processes and people



10%

Algorithms

- Provide employees with world-class AI models and proprietary ML capabilities (e.g., driver trees, rules engines, and context intelligence layers)

20%

Technology

- Seamlessly integrate AI tools into workflows
- Build a robust and scalable cloud infrastructure
- Provide employees with an AI and GenAI native workspace environment

70%

Processes...

- Drive vision and strategic clarity for an AI-enabled finance organization
- Design functional workflows of the future
- Build a transformation office to define and manage AI initiatives
- Define and monitor success metrics

...and people

- Have leadership set AI goals and identify AI champions to lead initiatives
- Assign OKRs across the organization, from functional leads down to individual practitioners
- Define and launch change management strategy, focusing on why AI works
- Assess change in skills required, and reskill employees based on needs

The time to act is now - to avoid falling behind in your finance AI journey

The technology is here

AI in finance is no longer experimental; most companies are already deploying AI across finance workflows

The value potential is rising

The increasing speed of AI adoption is widening the gap between leaders and laggards

First movers are reinventing capabilities

The role of finance is evolving from producers of numbers to architects of value

Talent expectations have shifted

Finance professionals expect AI enablement so they can focus on insight and value-adding tasks

Guardrails are critical


Ungoverned, self-directed AI usage can expose financial statements to audit and regulatory risk




Strategic clarity on North Star, with practical application of AI, drives transformative impact

BCG experts | Key contacts for AI in a finance transformation


North America



BCG X and Platinion



Europe, the Middle East, South America, and Africa



Asia-Pacific





BCG