THE FUTURE-READY CFO MASTERCLASS

MODULE 2 OF 20 | PILLAR I: AI, TECHNOLOGY AND DIGITAL INTELLIGENCE

Intelligent Finance Transformation

MODULE SUMMARY

This module shows how a finance function moves from a slow monthly close toward continuous close, predictive forecasting, and assisted controls. It explains a five step maturity path, the controls that must travel with automation, and the difference between process mining, robotic process automation, and AI agents. Worked examples compute close days saved, forecast error and bias, an anomaly test, and a three wave business case with payback. A fictional case, sharp boardroom questions, a practical roadmap template for the executive artifact, cited research from leading firms and standard setters, and and a ten question self assessment complete the module.

MODULE AT A GLANCE

Core question	How does a CFO move finance toward a continuous close, predictive planning, and assisted controls, in a sequence the organization can absorb, while keeping every control that a regulator, auditor, or director relies upon?
Why it matters to the CFO	Boards now expect finance to deliver speed and insight together. A transformation that saves days but weakens control, or that promises savings it cannot measure, damages the credibility of the person who sponsored it.
Key sources	COSO guidance on generative AI (through Deloitte, RSM US, EisnerAmper), PCAOB staff spotlight on generative AI, Deloitte CFO Signals (through the Journal of Accountancy), KPMG Belgium, PwC, and Deloitte close material.
Executive artifact	A finance transformation roadmap that names each wave, its owner, the controls that travel with it, the cost and benefit, and the success measures the board will see.

Important notice. This material is provided for educational and informational purposes only and does not constitute accounting, tax, legal, investment, or other professional advice. Readers should consult qualified professionals in accounting, tax, legal, and other relevant services before acting on any information presented. Fictional companies and figures are used for teaching and are labeled as such. The maturity ratings are an educational framework and have not been validated against a broad population of finance executives.

1. Executive overview and learning objectives

The core question. How does a finance function move from a monthly close that consumes the first two weeks of every month toward a function that closes continuously, forecasts with discipline, and monitors its own controls, and how does a chief financial officer do so without losing the reliability that the function exists to provide? The question is practical and not futuristic. A close that finishes in eight days instead of fifteen is a measurable result of work that many companies have already done, and the sequence of that work is the subject of this module.

Why this matters over the next five years. Deloitte reported in its fourth quarter 2025 CFO Signals survey, of 200 finance chiefs at North American businesses with at least one billion dollars of annual revenue, that 87 percent expect artificial intelligence to be very or extremely important to finance operations in 2026, and that 54 percent place the integration of AI agents among their top three finance transformation priorities (Journal of Accountancy, January 13, 2026). Enthusiasm of that kind invites hurried spending. The CFO who sponsors a transformation will be judged not by the quality of the tools purchased but by whether the close is faster, the forecast is closer to the outcome, the controls still operate, and the savings appear in the budget. Module 1 explains what the technology can and cannot do. This module explains how to apply it to the work of finance in a sequence that a finance team, an audit committee, and an external auditor can follow.

A word on terms before we proceed. **Intelligent finance** is a function in which routine processing is performed by software under rules and monitoring, so that people spend their time on judgment, exceptions, and advice. **Autonomous finance** is the far end of that path, a state in which software completes a process from start to finish with human review only on exceptions. The word is a direction and not a destination, since few controllers will allow software to close the books without a person reading the result. The practical goal for most companies in the next five years is a function that is largely touchless on routine items and well governed on the rest.

Learning objectives. After completing this module the reader should be able to do the following. Each objective is tested by the self assessment in Section 8.

1. Describe the five stages of a finance maturity path and place a function at its present stage, using evidence and not aspiration.
2. Explain how a continuous close differs from a faster monthly close, and compute the days saved by a change in the close calendar using the critical path.
3. Measure forecast quality with mean absolute percentage error and bias, compare a model with a naive benchmark, and explain why bias matters more to a board than a single large miss.
4. Specify the controls that must travel with automation, including segregation of duties, access, change control, logging, and human review thresholds.
5. Distinguish process mining, robotic process automation, and AI agents, and select the right tool for a given process.
6. Build a wave based roadmap with a business case, a payback period, a downside case, success measures, and named owners, and present it to a board.

How the module connects to the others. Module 1 supplies the vocabulary of generative and agentic systems that this module applies. Module 3 covers the data foundations without which none of the automation here can be trusted. Module 11 treats assurance over automated and AI assisted reporting, which the controls in this module must satisfy. Module 14 addresses the workforce that remains after routine work is automated, and Module 15 addresses the change leadership that decides whether a transformation lands. Technology cost

and return discipline appear in Module 4.

2. Conceptual foundations and frameworks

The maturity path. A transformation is easier to govern when it is described as a path with stages, since each stage has a different purpose, a different risk, and a different measure of success. The path in Exhibit 1 is a teaching framework developed for this module from common practice and is not a published standard. Its main lesson is that a function cannot skip a stage. A company that tries to reach predictive planning before it has one trusted set of data will produce confident forecasts from unreliable inputs, and a company that automates a process it has not first standardized will automate its inconsistencies.

Exhibit 1 A Five Stage Maturity Path for the Finance Function

Stage	What the function does	Typical evidence	Principal risk
1. Manual and fragmented	Spreadsheets, email approvals, local reconciliations, and a close that depends on a few experienced people	Close of fifteen or more business days. Many manual journal entries. Forecast rebuilt by hand each cycle	Key person dependence and errors that surface late
2. Standardized	One chart of accounts, one close calendar, written procedures, and shared service teams for routine work	Close checklist with owners. Reconciliations on a fixed schedule. Defined data owners (Module 3)	Standard processes that are slow because nothing is automated
3. Automated	Rules based automation for matching, accruals, intercompany, and reporting. Exceptions are routed to people	Touchless rate measured. Reconciliations by exception. Close shortened by several days	Automation that conceals errors if exceptions are not reviewed
4. Predictive and assisted	Driver based plans, rolling forecasts, statistical models, and monitoring of controls and transactions	Forecast error tracked against a naive benchmark. Continuous controls monitoring operating	Models trusted beyond their evidence. Alerts that no one investigates
5. Autonomous on routine	Software completes defined processes end to end. People set policy, handle exceptions, and judge the outcomes	Soft close each day. Human review by threshold. Full audit trail and logged approvals	Oversight that thins under efficiency pressure

Source: teaching framework prepared for this module from common practice. It is not a published standard and has not been validated empirically. The COSO point on thinning oversight is from RSM US (April 7, 2026).

Continuous close. The traditional close treats the month as a batch. Transactions accumulate for thirty days, and then a team spends days reconciling, accruing, matching, and reviewing, all while the books are frozen. A **continuous close** spreads that work across the month, so that on the last day little remains to be done. Subledgers are matched daily, bank activity is reconciled as it arrives, intercompany balances are agreed as they arise, and accruals are estimated from live data and refreshed. The month end becomes a review of a ledger that is already nearly final, which some practitioners call a soft close, and the days saved come from removing work from the critical path and not from asking people to type faster. A Deloitte overview of next generation close material, which carries a 2021 copyright notice, describes the aim as a touchless and continuous close, and states vendor expectations of a 40 to 50 percent reduction in time to close and a best in class close of one to two business days. These are claims of the vendor and not independent findings, and a reader should treat them as aspiration to test and not as a promise.

The critical path. The tool that explains why a close takes as long as it does is the critical path, which is the longest chain of tasks that each depend upon the one before. Automating a task that is not on the critical path

saves effort but saves no calendar days. In the worked example in Section 3 the bank reconciliation takes three days, yet the close cannot finish sooner by shortening it alone, because the accrual task takes the same three days and both must finish before the balance sheet reconciliations begin. A finance leader who funds the wrong task will spend money and report no improvement, and the error is common.

Intelligent planning. Planning has three layers that the function should keep distinct. **Driver based planning** builds the budget from the operating quantities that cause revenue and cost, such as units, price, hires, and utilization, so that a change in a driver flows through the plan without a rebuild. A **rolling forecast** extends the forecast horizon each period, for example always showing the next four or six quarters, so that management looks forward at a constant distance and not toward a year end that draws closer each month. **Predictive forecasting** uses statistical or machine learning models to project a line, such as revenue or cash receipts, from history and drivers, and the model's value is measured by whether it beats a simple benchmark. The three layers work together. Drivers explain the plan, the rolling horizon keeps the plan current, and models supply the first estimate that people then challenge.

Exhibit 2 Four Layers of Intelligent Planning

Planning layer	Question it answers	Typical tools	Measure of success
Driver based plan	What operating assumptions produce the numbers, and what changes if they move?	Planning platform linked to operating data. Named driver owners in sales, operations, and people teams	Share of plan lines tied to a driver. Cycle time to re-plan after a change
Rolling forecast	Where do we expect to land over the next four to six quarters?	Quarterly or monthly refresh. Short template. Variance bridge to the prior forecast	Days to publish. Forecast error by horizon. Management use in decisions
Predictive model	What does history and the driver data suggest before judgment is added?	Time series methods, regression, and machine learning on clean history (Module 3)	Error against a naive benchmark. Bias. Stability across periods
Scenario generation	What happens under a different set of assumptions, and how likely is each?	Driver sets for base, downside, and upside. Software that proposes scenarios from a prompt, reviewed by a person	Decision made earlier. Scenarios used by the board (Module 17)

Source: author's synthesis of common practice. KPMG Belgium (September 2025) reports that Belgian CFOs are working on rolling forecasts and predictive analytics as part of agile finance processes, which supports the direction but not the table's specifics.

Scenario generation. Generative systems can propose scenarios quickly from a prompt, for example by drafting a downside case for a supplier failure or a rate increase, and can document the assumptions. The speed is useful and the hazard is plain. A scenario is only as good as the driver relationships behind it, and a language model does not know the company's contracts or constraints unless those are provided. The CFO should treat generated scenarios as a first draft that an analyst tests against the driver model, and should require that every scenario show its assumptions in a table that a director can read.

Assisted controls. A control is an activity that reduces the chance that an error or fraud reaches the financial statements. Traditional controls are sampled, periodic, and often performed after the fact. **Continuous controls monitoring** tests all transactions, or a very large share of them, against rules on a continuing basis and raises exceptions. **Anomaly detection** goes further, by learning what a normal journal entry or payment looks like and flagging those that differ, for example an entry posted at an unusual hour or a payment just under an approval limit. The software is an additional reviewer and not a replacement for the owner of the control. The PCAOB staff, in a July 2024 spotlight, noted that preparers who use generative AI cited risks of data privacy, unreliable output, auditability of source data and generated content, and amplification of existing

information technology risks such as segregation of duties.

The COSO lens. In February 2026 the Committee of Sponsoring Organizations of the Treadway Commission, known as COSO, published guidance titled *Achieving Effective Internal Control Over Generative AI*. Deloitte's April 2026 summary says that it builds on the 2013 Internal Control Integrated Framework, maps risks to the framework's 17 principles, and recommends a six step roadmap of govern, inventory, assess, design, implement, and monitor. Two ideas in the guidance matter for this module. The first is that generative systems are probabilistic, which means that the same input can produce different outputs, and so a one time approval is not enough and monitoring must continue. The second is that human involvement should be scaled to the risk of the use case and the degree to which its output drives decisions. RSM US adds that once a process would fail or change materially without the AI, the AI should be governed as a control.

Process mining, robotic process automation, and AI agents. Three tools are often confused. **Process mining** reads the event logs that systems already keep and reconstructs how a process actually runs, showing every path, delay, and rework loop. It is a diagnostic. **Robotic process automation**, or RPA, is software that imitates a person following a fixed sequence of clicks and entries in an application. It is deterministic, cheap to start, and fragile when a screen changes. An **AI agent** is a system that pursues a goal by choosing among tools and steps, often with a language model reading documents and deciding what to do next. It handles variation that rules cannot, and it is probabilistic. Exhibit 3 compares them. The sequence that works in practice is to mine first, standardize second, automate with rules where the rules are stable, and use agents only where judgment about unstructured material is needed and review is built in.

Exhibit 3 Process Mining, Robotic Process Automation, Rules and AI Agents Compared

Tool	What it does	Best suited to	Weakness	Control focus
Process mining	Reconstructs actual process flow from system event logs and measures delay and rework	Finding where the close and procure to pay actually lose time, before spending on automation	Needs clean event data. Shows what happens and not why	Data access, completeness of logs, privacy of employee activity
Robotic process automation	Repeats a fixed series of keystrokes and entries in existing applications	High volume, stable, rule based tasks such as copying data between systems	Breaks when screens or rules change. Can automate a bad process	Bot credentials, change control, logging, owner of each bot
Rules and workflow automation	Matches, routes, and approves items by explicit rules inside the ERP or a platform	Three way match, bank reconciliation, intercompany matching, approvals	Rules age as the business changes	Rule change approval and periodic rule review
AI agent	Pursues a goal across tools. Reads documents, drafts entries, proposes actions	Variable inputs such as invoices, contracts, emails, and exception analysis	Probabilistic. Can be wrong with confidence. Hard to explain	Review threshold, tool permissions, prompt and model change control, logging (COSO, Module 10)

Source: author's synthesis. The COSO capability view and the point that deterministic automation may be safer than generative AI where it suffices are from EisnerAmper (May 13, 2026) and Deloitte DART (April 3, 2026).

Autonomous finance in context. It is helpful to observe what vendors and firms are now saying. PwC announced in May 2025 an agent platform for operations and finance leaders that it describes as able to detect leakage and act in systems such as ERP, and the announcement gave no statistics. The KPMG Belgium event summary of September 2025, covering a survey of 157 Belgium based CFOs, reported that many

organizations remained at an early stage of AI use, mainly through RPA, and that 66 percent had difficulty finding people with the right skills. The two views are consistent. The offering is advancing faster than the average finance function has been able to prepare its data, controls, and people to absorb it.

3. Practical applications and tools

Step one: establish the baseline. No transformation can be judged without a baseline measured before the work begins. For each major process, record the cycle time, the touchless rate, the error rate, the hours of effort, the number of manual entries, and the cost per transaction. Process mining, where data allows, gives a fact based view, and where it does not a two week manual sample will do. Record the baseline in a table that the audit committee can see, because a benefit that was never measured at the start cannot be proven at the end.

Step two: map the close calendar and find the critical path. Harbor Ridge Industrial, Inc. is a fictional manufacturer with several entities. Its close takes 15 business days. Exhibit 4 lists nine close tasks, their duration in business days before and after a first wave of automation, and the finish day of each computed from the task dependencies. The calculation was run in Python. A task starts when the last task it depends upon has finished, and the close ends when the reporting package is issued.

Exhibit 4 Harbor Ridge Industrial, Inc. (Fictional): Close Calendar Before and After Wave Two

Close task	Depends on	Before: days	Before: finish day	After: days	After: finish day
Subledger cutoff and posting	None	2	2	1	1
Bank and cash reconciliation	cut	3	5	1	2
Intercompany matching and elimination	cut	4	6	1	2
Accrual preparation	cut	3	5	1	2
Balance sheet reconciliations	accr, cash	4	9	2	4
Journal entry review and posting	accr, ic	2	8	1	3
Consolidation and currency translation	bs, je, ic	2	11	1	5
Management review and variance commentary	con	3	14	2	7
Reporting package issued	rev	1	15	1	8
Close completed (business days)			**15**		**8**

Source: fictional example, computed in Python. Dependency keys: cut is subledger cutoff, cash is bank reconciliation, ic is intercompany, accr is accruals, bs is balance sheet reconciliations, je is journal entries, con is consolidation, rev is management review. Days are business days.

Reading the calendar. Before the change, the critical path ran through cutoff, accruals, balance sheet reconciliations, consolidation, review, and the package, for 15 days. The intercompany task took four days, the longest single task other than the balance sheet reconciliations, yet it was not on the critical path, because it finished on day 6 while the balance sheet reconciliations did not finish until day 9. After automation, the close fell to 8 business days, a saving of 7 days or 46.7 percent. The critical path was the same chain of tasks as before, only shorter at each step, and this is a useful check: the days saved came from shortening the tasks that were on the path. A third column of the model, which represents a continuous close with daily matching and an exceptions only review, produces a close of 4.5 business days, a saving of 10.5 days or 70.0 percent

against the starting point. Moving from the second state to the third is far harder than moving from the first to the second, because it requires data quality and soft close discipline in every entity.

The effort model. Days saved on the calendar are not the same as hours saved. Preparer effort at Harbor Ridge fell from 1,180 hours to 640 hours per close, a saving of 540 hours. Over 12 closes this is 6,480 hours, which at a loaded cost of 62 dollars per hour is 401,760 dollars a year. The CFO should note what this figure is. It is capacity released, and it becomes a saving only if the hours are redeployed to work of higher value or the headcount is reduced through attrition, a choice discussed in Module 14.

Controls that travel with the automation. The test of a sound transformation is that every control that existed before still operates after, in a form that an auditor can test, and that new risks created by the automation have their own controls. Exhibit 5 sets out the controls that should be designed into each automation. Many of them follow the building blocks that COSO recommends for generative systems, which Deloitte lists as access and acceptable use restrictions, input and data controls, prompt and configuration change control, output validation and exception handling, logging and traceability, and monitoring for drift, anomalies, and unauthorized use.

Exhibit 5 Controls That Must Travel with Automation

Control area	What it requires	Evidence an auditor will ask for	Owner
Access and permissions	Each bot, agent, and user has the least access needed. Service accounts are named and reviewed each quarter	Access listing, approval of each grant, quarterly review sign off	Controller and IT security
Segregation of duties	No single user, bot, or agent may prepare, approve, and release the same item. Automation does not remove the conflict	Conflict matrix, test of conflicting roles, exceptions log	Controller
Input and data controls	Source data are complete, authorized, and from approved systems. Changes to feeds are controlled	Reconciliation of input counts and amounts to the source system	Data owner (Module 3)
Change control	Changes to rules, prompts, models, and configuration are requested, tested, approved, and recorded	Change tickets with test evidence and approver	Process owner and IT
Output validation and exceptions	Outputs are tested against expected ranges. Items outside a threshold go to a person	Exception queue, aging of open items, sign off	Process owner
Human review threshold	A written rule states which items a person must review, based on amount, risk, and error history	Policy, sampling evidence, override percentage	CFO and Controller
Logging and audit trail	Every automated action records who or what acted, when, on what data, and with what result. Logs cannot be altered by the actor	Log extract for a sample, retention policy	IT and Internal Audit
Monitoring for drift	Volumes, amounts, override rates, and error rates are tracked so that a change in behavior is noticed	Monitoring report and follow up on thresholds breached	Process owner and Internal Audit

Source: author's synthesis, built on the COSO building blocks summarized by Deloitte DART (April 3, 2026), the RSM US article of April 7, 2026 on evidence and reliance thresholds, and the PCAOB staff spotlight of July 2024. Module 11 discusses assurance over these controls.

Step three: measure forecast quality honestly. The worked example uses twelve months of fictional revenue for Harbor Ridge, in thousands of dollars, with two forecasts made one month ahead: a manual forecast built by the business units and a model forecast built from drivers and history. Two measures are used. The **mean absolute percentage error**, or MAPE, is the average of the absolute percentage miss each month, and it measures how far the forecast typically is from the outcome. **Bias** is the average of the signed percentage errors, positive when the forecast is above the outcome, and it measures whether the forecast leans one way. A forecast with zero bias and a large MAPE is noisy but honest, while a forecast with a small MAPE and a persistent bias tells the board something is wrong with the incentives. A third measure, the **weighted absolute percentage error**, or WAPE, is the total absolute miss divided by total actual, and it prevents small months from distorting the picture.

Exhibit 6 Harbor Ridge Industrial, Inc. (Fictional): Monthly Revenue Forecast Accuracy, Thousands of Dollars

Month	Actual	Manual forecast	Manual error %	Model forecast	Model error %
Jan	4,200	4,400	+4.8	4,180	-0.5
Feb	4,310	4,500	+4.4	4,350	+0.9
Mar	4,450	4,600	+3.4	4,420	-0.7
Apr	4,380	4,550	+3.9	4,410	+0.7
May	4,520	4,650	+2.9	4,560	+0.9
Jun	4,690	4,800	+2.3	4,650	-0.9
Jul	4,610	4,780	+3.7	4,650	+0.9
Aug	4,740	4,900	+3.4	4,700	-0.8
Sep	4,880	5,000	+2.5	4,850	-0.6
Oct	4,820	5,050	+4.8	4,850	+0.6
Nov	5,010	5,150	+2.8	4,980	-0.6
Dec	5,150	5,300	+2.9	5,120	-0.6
MAPE (percent)			**3.47**		**0.72**
Bias (percent)			**+3.47**		**-0.05**
WAPE (percent)			**3.44**		**0.72**

Source: fictional data, computed in Python. Error is (forecast minus actual) divided by actual. A positive error means the forecast was too high.

Reading the results. The manual forecast had a MAPE of 3.47 percent and a bias of +3.47 percent, and it was above the outcome in 12 of the twelve months. The model forecast had a MAPE of 0.72 percent and a bias of -0.05 percent, being above the outcome in 5 months. The consistent overstatement in the manual forecast is the more important finding than its size. A business that always forecasts high is signaling a target culture and not an estimate, and a board that sees this pattern will discount every forecast it receives.

The naive benchmark. A model proves its worth only against a simple alternative. The **naive forecast** predicts that next month will equal this month. Over the eleven months from February to December, the naive forecast had a MAPE of 2.65 percent, the manual forecast 3.35 percent, and the model 0.74 percent. The model beat the naive benchmark, which is the minimum standard, and the manual forecast did not. Because revenue in this example trends upward, the naive method lags the trend in every month, and a CFO who tests a vendor

model should request this comparison on the company's own history before signing a contract.

Step four: apply assisted controls to payments and entries. Exhibit 7 shows a small test on twelve fictional payments against an approval limit of 10,000 dollars. Five simple rules were applied: an amount from 90 percent up to the limit, a repeat of the same vendor and amount within five days, a weekend posting date, a preparer who is also the approver, and a round amount of 5,000 dollars or more. The rules flagged 6 of the twelve payments. Investigation found that two were genuine problems, a duplicate payment to a vendor and a self approved payment, so precision, the share of flags that were real, was 33.3 percent, and recall, the share of real problems caught, was 100.0 percent.

Exhibit 7 Twelve Fictional Payments Tested Against Five Rules (Approval Limit 10,000)

ID	Vendor	Amount	Date	Preparer	Approver	Flags raised
P01	A	4,820.00	Tue Sep 1	Kim	Lee	None
P02	B	9,950.00	Wed Sep 2	Ray	Lee	near limit
P03	B	9,950.00	Fri Sep 4	Ray	Lee	near limit, duplicate
P04	C	7,000.00	Sat Sep 5	Ray	Ray	weekend, self approval, round amount
P05	D	1,284.37	Mon Sep 7	Kim	Lee	None
P06	E	15,200.00	Thu Sep 10	Kim	Ortiz	None
P07	F	5,000.00	Tue Sep 8	Kim	Lee	round amount
P08	A	3,310.55	Wed Sep 9	Kim	Lee	None
P09	G	9,400.00	Thu Sep 10	Ray	Lee	near limit
P10	H	640.10	Fri Sep 11	Kim	Lee	None
P11	I	2,250.00	Sun Sep 13	Ray	Lee	weekend
P12	J	8,075.20	Mon Sep 14	Kim	Lee	None

Source: fictional data, rules applied in Python. The rules are illustrations and not a complete fraud program. Weekend flags and round amounts are weak signals on their own.

What the test teaches. Four of the six flags were false alarms, and that is typical of a first rule set. A flag is a prompt to look and not a finding. The practical lessons are three. First, combinations are stronger than single rules, since the two genuine problems each carried either two or three flags. Second, every flag must be closed with a recorded disposition, so that the false alarm rate can be measured and the rules tuned. Third, an alert that no one investigates is worse than no alert, because it creates evidence that the company knew and did not act. Anomaly models learn patterns from history and can find what rules miss, but they raise the same need for dispositions and for a person who owns the queue.

Segregation of duties in automated workflows. Segregation of duties means that no one person controls all the steps of a transaction, such as creating a vendor, entering an invoice, and approving payment. Automation tends to concentrate these steps, because a single bot or agent can be given broad access for convenience. The remedy is to treat each bot and each agent as a user with a role, to include it in the conflict matrix, and to forbid the same automated identity from both preparing and releasing an item. Where an agent prepares an entry, a different person or a separately controlled process must approve it above the threshold, and the log must show both.

Step five: sequence the work in waves and test the business case. A wave is a bounded set of changes with its own owner, controls, cost, and success measure, and the next wave begins only when the gate for the present wave is passed. For Harbor Ridge, the three waves are shown in Exhibit 8 with their costs and benefits in thousands of dollars. Benefits are shown at full run rate and are realized gradually, because a new process takes quarters to reach its planned performance.

Exhibit 8 Harbor Ridge Industrial, Inc. (Fictional): Three Wave Program, Thousands of Dollars

Wave	Scope	One time cost	Annual run cost	Annual benefit at full run rate	Gate to proceed
Wave 1: Foundation and quick wins	Close calendar standardized, process mining of close and payables, bank and intercompany matching, data owners named	900	180	700	Baseline measured. Touchless rate and close days improving
Wave 2: Automate the close and controls	Accrual and reconciliation automation by exception, continuous controls monitoring, segregation of duties redesign	1,600	320	1,500	Parallel run clean. Auditor agrees on evidence
Wave 3: Predictive and continuous	Driver based plan, rolling forecast, predictive models, soft close each day, agent assisted analysis	1,300	280	1,300	Forecast beats the naive benchmark. Bias within tolerance
Total		3,800	780	3,500	

Source: fictional example. Totals computed in Python.

Exhibit 9 Harbor Ridge Industrial, Inc. (Fictional): Net Cash Flow and Payback

Thousands of dollars	Year 1	Year 2	Year 3	Year 4
Net cash flow, base case	(1,092)	(626)	1,006	2,720
Cumulative, base case	(1,092)	(1,718)	(712)	2,008
Net cash flow, benefits at 70 percent	(1,176)	(971)	229	1,670
Cumulative, benefits at 70 percent	(1,176)	(2,147)	(1,918)	(248)

Source: fictional example, computed in Python. Net cash flow equals realized benefit less realized run cost less one time cost. Realization by wave: Wave 1 at 40, 100, 100, 100 percent; Wave 2 at 0, 30, 100, 100; Wave 3 at 0, 0, 30, 100.

Reading the business case. The base case requires 3,800 thousand dollars of one time cost, produces cumulative cash of (1,718) thousand dollars at the low point in year two, and pays back in 3.26 years. The net present value at a 10 percent discount rate, over four years and with no value assigned beyond year four, is 1,104 thousand dollars. When benefits are reduced to 70 percent of plan, the net present value falls to (559) thousand dollars and the payback is not reached within four years. The CFO should therefore ask for the downside case before approving the program, because the answer shows how much of the case rests on benefits that are hard to prove. Finance leaders sometimes include in the benefit figure the full value of released hours. A cautious approach counts only those hours that were removed from the budget or redeployed to a measured task, and states the rest as capacity.

Step six: measure that finance is truly faster and better. A transformation succeeds when the measures move and stay moved. Exhibit 10 lists the measures a board should see each quarter, with the Harbor Ridge

baseline and the target. The touchless rate is the share of transactions processed without human touch, and the cost to serve is the finance cost of supporting a business activity such as a processed invoice.

Exhibit 10 Measures That Show Finance Is Faster and Better

Measure	Definition	Baseline	Year one result	Why it matters
Close cycle time	Business days from period end to issued package	15	8	Speed of decision information
Touchless invoice rate	Invoices processed with no human touch divided by all invoices	41.0 percent	78.0 percent	Real automation and not partial automation
Cost per invoice	Total payables cost divided by invoices processed	$9.20	$5.10	Cost to serve and its trend
Forecast MAPE and bias	Average absolute and signed percentage error, by horizon	3.47 and +3.47	0.72 and -0.05	Quality and honesty of planning
Manual journal entries	Share of entries keyed by a person and not generated	Measure at baseline	Falling, with reasons for each remaining entry	Residual risk of error and override
Control exceptions closed	Alerts closed with a recorded disposition within a set time	Not measured	At least 95 percent in the target period	Proof that monitoring leads to action
Headcount redeployment	People moved to analysis, controls, or business partnering, with the new role stated	Not measured	Named roles and hours	Whether capacity became value (Module 14)

Source: fictional baseline and results for Harbor Ridge, computed in Python (touchless rate 39,360 of 96,000 invoices before and 74,880 after; payables cost 883,200 and 489,600 dollars). The 95 percent target is a proposed example and not a standard.

Common change risks. Five risks recur. The first is automating before standardizing, so that each entity keeps its own variations and the automation is built several times. The second is neglecting data quality (Module 3), which makes every model and every rule unreliable. The third is thin oversight, in which a review that began as thorough becomes a click through as volumes rise. The fourth is the loss of knowledge, because the experienced people who knew why a reconciliation was done in a certain way leave, and the reasons are not written down. The fifth is cultural, in that a team which fears that automation is a plan for dismissal will conceal the very process facts that the program needs. Module 15 covers the leadership of these risks.

4. Executive case study and decisions

Fictional case: Alder Peak Components, Inc. The names, companies, and figures in this case are fictional and exist only for teaching. Alder Peak Components makes engineered parts, with annual revenue of 620 million dollars, fourteen legal entities in six countries, three ERP instances inherited from acquisitions, and a monthly close of twelve business days. The company is owned by public shareholders and the audit committee is attentive to internal control over financial reporting. Its chief executive has told the board that the close should reach five business days within twelve months, and a software vendor has proposed a single platform with AI agents that would post accruals, reconcile accounts, and prepare consolidation entries with, in the vendor's words, minimal human involvement.

The decision. The CFO must decide whether to commit to the vendor proposal, to adopt a staged program, or to buy tools without redesigning the process. The board meets in three weeks. Three facts complicate the

choice. The three ERP instances use different charts of accounts for four of the entities. The external auditor has said that it will need evidence of how any automated posting is controlled and that it will test the controls in the first year of operation. And the finance team of 46 people includes several long serving accountants whose knowledge of the intercompany arrangements is not documented.

The options. Exhibit 11 compares three options over five years, in millions of dollars, with a net present value at 10 percent and a second value in which benefits are 70 percent of plan. The figures are fictional and the calculations were run in Python.

Exhibit 11 Alder Peak Components, Inc. (Fictional): Options Compared, Millions of Dollars

Option	One time cost	Year 1 net	Year 5 net	Payback (years)	NPV at 10 percent	NPV, benefits at 70 percent
A. Single platform, autonomous posting in 12 months	11.0	(6.8)	5.1	3.82	2.23	(3.17)
B. Three waves with decision gates	8.4	(2.6)	5.0	3.26	5.08	0.24
C. Tools only, no process change	1.4	(1.0)	1.2	1.83	2.55	0.95

Source: fictional example, computed in Python. Option A: one time cost 6.0 and 5.0 in years one and two, benefits 0, 4.4, 7.0, 7.0, 7.0. Option B: one time 3.2, 3.0, 2.2, benefits 0.9, 3.2, 5.6, 6.5, 6.5. Option C: one time 1.4, benefits 0.8 then 1.6. Run costs are included in net figures.

Reading the table. Option A has the largest benefit in the later years and the most attractive headline, but it requires 11.0 million dollars up front, pays back in 3.82 years, and has a negative net present value of (3.17) million dollars if benefits are 30 percent below plan. Option B has the highest net present value in the base case, 5.08 million dollars, pays back in 3.26 years, and remains positive at 0.24 million dollars in the downside. Option C is safe and small, with a payback of 1.83 years, a net present value of 2.55 million dollars, and a downside value of 0.95 million dollars, which is better than Option B in the downside. It leaves the process as slow as before, since it buys tools without changing the work, and it cannot approach the close target the board has set.

The recommended path. The CFO should recommend Option B and should tell the board plainly that it will not reach a five day close in twelve months, and that it can reach about eight days in that period and the five day target in about twenty four months. The reasoning has four parts. First, the critical path at Alder Peak runs through accruals, balance sheet reconciliations, and consolidation across three ERP instances, so a platform cannot shorten it until the charts of accounts are mapped and intercompany matching is automated. That work belongs to the first wave. Second, the auditor will test the controls over automated posting, and a program that goes live in twelve months leaves little time for a parallel run. Third, the downside case shows that a single platform commitment is the only option that loses value if benefits fall short. Fourth, a staged program lets the board see evidence at each gate, and Module 9 reminds directors that oversight is strongest when evidence arrives before the next commitment and not after.

Decision gates the CFO should propose. Autonomous posting of accruals should begin only for items below a set amount, for example 25,000 dollars, after a parallel run of at least three closes in which the automated entries differ from the manual entries by less than a tolerance that the controller and auditor have agreed in writing. The threshold should rise only on evidence. Every automated entry should carry the identity of the process that created it and the identity of the person who approved the rule, and any entry above the threshold should require a human approver. The program should also report to the audit committee each quarter on the measures in Exhibit 10, on exceptions closed, and on any control deficiency, with Internal Audit performing an independent review of the first wave.

Handling the people risk. The CFO should name the long serving accountants as the authors of the new process documentation and should say so openly, with a statement of what roles they will hold after automation. Module 14 discusses redeployment, and Module 15 discusses how to retain the people whose knowledge the program needs. The savings in Exhibit 11 do not depend on dismissals, and the CFO should resist the temptation to promise them, since a promise of headcount reduction made early can cost the program the cooperation on which its schedule depends.

What a weaker answer misses. A weaker answer accepts the twelve month target because the chief executive stated it, and signs the vendor proposal on the strength of its demonstration. It reports savings in hours and not in dollars, and it does not run a downside case. It lets the vendor define how the agents will be controlled, so that the auditor first learns of the design in the fourth quarter. It treats the intercompany knowledge as an afterthought, and it commits the full budget on the first day. Each of these omissions is common, and each is avoided by the discipline of waves with gates.

Questions for reflection. What would you tell the chief executive if he insisted on the twelve month target? Which single control would you place at the first gate, and why? If the auditor refused to rely on the automated accrual process in the first year, how would the business case change, and would you still proceed? Write your answers before reading the boardroom questions in Section 5.

5. Boardroom discussion and key takeaways

Questions directors should ask, and the CFO should be ready to answer. The questions below follow the order of this module. A director need not understand the technology to ask them, and a CFO who cannot answer them is not ready to ask for the investment.

1. What is the baseline today for close days, touchless rate, forecast error, and cost per transaction, and who measured it? A program that cannot show its starting point cannot show its result.
2. Which tasks are on the critical path of the close, and does the spending address those tasks first?
3. Which processes are standardized and which are not, and what happens if we automate a process that varies by entity?
4. Where does the program use rules, where does it use AI agents, and why? What is the review threshold for each agent, and what is the evidence that it has been followed?
5. How are bots and agents included in the segregation of duties matrix, and who can see the logs?
6. How does the forecast model compare with a naive forecast and with the manual forecast, and what is the bias? Has bias fallen since the program started?
7. What is the payback if benefits arrive at 70 percent of plan, and what would cause that outcome?
8. What has the external auditor said it needs in order to rely on the automated controls, and when was the auditor consulted?
9. What becomes of the people whose work is automated, and how will we know that released capacity became something of value?

Key takeaways. The ten lessons that follow condense the module.

1. Transformation is a sequence of stages, and a function cannot skip one. Standardize before automating, and clean the data before building models.
2. A continuous close removes work from the critical path. Speeding up a task that is off the path saves effort and no calendar time.

3. Days saved and hours saved are different results. Hours become a saving only when they are removed from the budget or redeployed to measured work.
4. Forecast quality has two parts, accuracy and bias. A model must beat the naive benchmark on the company's own history.
5. Controls travel with automation. Every control that existed before must operate after, and each bot or agent needs its own access, logging, and review threshold.
6. Segregation of duties applies to software. A single automated identity should never prepare and release the same item.
7. An alert without a recorded disposition is a liability. Measure false alarms and tune the rules.
8. Choose the tool for the process: mine to learn, rules for stable work, RPA for narrow gaps, and agents only where judgment about unstructured material is required and review is built in.
9. Test the business case against a downside, and prefer the program that stays positive when benefits disappoint.
10. Report progress to the board with measures, and let each wave earn the next through a gate.

A closing thought for the section. The temptation in every transformation is to describe the destination in the language of the vendor and the journey in the language of hope. The more reliable method is the plain one that a careful investor applies to any capital project: measure the starting point, state the expected return, test the downside, and check the result when the money has been spent. Finance has the unusual advantage that its own work is among the easiest in the company to measure, and a CFO who applies the discipline to the finance function earns the credibility to ask the same of everyone else.

6. The executive artifact: a finance transformation roadmap

The artifact for this module is a one page roadmap that a CFO can show to a board. It names each wave, its owner, the controls that travel with it, the cost and benefit, the gate, and the success measures. Exhibit 12 is a filled in example for Harbor Ridge, and the reader should complete a version for the company's own function. Each entry should be specific enough that a director can later check whether it was delivered. Costs are in thousands of dollars and are consistent with Exhibits 8 and 9.

Exhibit 12 Harbor Ridge Industrial, Inc. (Fictional): Finance Transformation Roadmap

Item	Wave 1: Foundation	Wave 2: Automate close and controls	Wave 3: Predictive and continuous
Timing	Months 1 to 9	Months 7 to 18	Months 16 to 30
Executive owner	Controller	Controller with Chief Information Officer	VP of FP&A with Controller
Outcome sought	Standard close calendar. Measured baseline. Quick wins in matching	Close of about eight days. Reconciliations by exception. Controls monitored	Close of about four to five days. Rolling forecast. Models beating the naive benchmark
Controls that travel	Access review, close checklist sign off, data owners named	Segregation of duties matrix with bots, change control, logging, review thresholds	Model change control, bias monitoring, scenario assumptions table, agent permissions

Item	Wave 1: Foundation	Wave 2: Automate close and controls	Wave 3: Predictive and continuous
One time cost and annual run cost	900 and 180	1,600 and 320	1,300 and 280
Annual benefit at full run rate	700	1,500	1,300
Success measures	Close days. Touchless rate. Hours per close	Close days. Manual entries. Exceptions closed. Cost per invoice	MAPE and bias. Days to re-plan. Decisions informed by the forecast
Gate to proceed	Baseline signed off by the audit committee chair	Clean parallel run. Auditor agrees on evidence	Forecast beats naive. Bias inside tolerance for two quarters
Principal risk and response	Automating variation. Respond by standardizing first	Thin oversight. Respond with review thresholds and Internal Audit review	Models trusted too far. Respond with human challenge and drift reports
Links to other modules	Module 3 (data), Module 15 (change)	Module 11 (assurance), Module 10 (risk)	Module 14 (workforce), Module 17 (scenarios)

Source: fictional example consistent with Exhibits 8 and 9. The reader should replace every entry with the company's own facts, owners, and dates.

How to complete the artifact. Begin with the baseline, since every other entry depends on it. Then name one executive owner for each wave, and avoid a committee as owner. Write each success measure so that a person who was not in the room could calculate it from data. Insert the gate before the cost, so that the first conversation about the next wave is about evidence. Finally, ask the external auditor and the head of Internal Audit to read the controls row, and record any change they request. The roadmap should be reviewed every quarter, and it should be treated as a living document and not a record of the original plan.

7. Research and Perspectives from the Leading Firms

The publications below were found through search and each page was opened before it was described. Where a page returned only part of its content, the description states only what was returned. Several concern the COSO guidance on generative AI, which is the most relevant recent standard setting development for assisted controls. Guidance changes quickly, and a reader should confirm the current edition before relying on it.

Exhibit 13 Publications of the Leading Firms and Standards Bodies

Source	Title and URL	What it says
Deloitte DART	COSO Releases Publication on Internal Controls Related to Generative AI, Heads Up, April 3, 2026. https://dart.deloitte.com/USDART/home/publications/deloitte/heads-up/2026/coso-internal-controls-generative-ai	Summarizes the COSO guidance of February 23, 2026, which builds on the 2013 framework and recommends a six step roadmap and control building blocks including logging, output validation, and change control
RSM US	Is generative AI being used to make decisions you cannot account for? (COSO aligns AI governance with internal control guidance), April 7, 2026. https://rsmus.com/insights/services/risk-fraud-cybersecurity/coso-aligns-ai-governance-with-internal-control-guidance.html	States that policies alone are insufficient, that controls must operate and leave evidence, and that AI which a process cannot do without should be governed as a control

Source	Title and URL	What it says
EisnerAmper	Article on COSO AI governance and the internal control framework, May 13, 2026 (the exact page title was not captured). https://www.eisneramper.com/insights/artificial-intelligence-insights/coso-ai-governance-internal-control-framework-0526/	Describes capability types of generative AI, the probabilistic nature of outputs, shadow AI, and the advice to check whether deterministic automation would be safer before using generative AI
PCAOB	Spotlight: Staff Update on Outreach Activities Related to the Integration of Generative Artificial Intelligence in Audits and Financial Reporting, July 2024. https://pcaobus.org/documents/generative-ai-spotlight.pdf	Reports staff outreach finding early stage use by auditors and preparers, and risks of data privacy, unreliable output, auditability, and amplified information technology risks
Deloitte (through the Journal of Accountancy)	CFOs signal crucial role that technology will play in 2026 (title taken from the web address), January 13, 2026. https://www.journalofaccountancy.com/news/2026/jan/cfos-signal-crucial-role-that-technology-will-play-in-2026	Reports Deloitte Q4 2025 CFO Signals survey of 200 CFOs: 87 percent expect AI to be very or extremely important to finance in 2026, and 54 percent rank integrating AI agents among top three transformation priorities
KPMG Belgium	CFOA Next Gen Finance: the CFO as strategic co-pilot in a digital and AI-driven world (title taken from the web address), September 2025. https://kpmg.com/be/en/home/insights/2025/09/cfoa-next-gen-finance-the-cfo-as-strategic-co-pilot-in-a-digital-and-ai-driven-world.html	Event summary with a survey of 157 Belgian CFOs: early stage AI use mainly through RPA, 62 percent prioritizing process standardization and technology transformation, and 66 percent reporting difficulty finding skills
PwC	Agent Powered Performance, press release, May 16, 2025. https://www.pwc.com/us/en/about-us/newsroom/agent-powered-performance.html	Announces an agent platform that acts in enterprise systems. Promotional in tone and gives no statistics, so its performance claims are not independently verified
Deloitte	Next Gen Close, Consolidation, and Reporting, solution overview, copyright 2021. https://www.deloitte.com/content/dam/assets-shared/legacy/docs/gx-next-gen-finance-assets-ccr.pdf	Describes a touchless, continuous close through automation, with stated expectations of 40 to 50 percent less time to close. These are the firm's own claims with no sample stated

Each page was opened and the descriptions follow what was returned. No EY, BDO, Grant Thornton, Baker Tilly, Crowe, Forvis Mazars, CBIZ, Moss Adams, CohnReznick or CLA publication was read for this module, and none is cited. The COSO publication itself and NIST material were not opened, and COSO content is cited only as summarized by Deloitte, RSM US and EisnerAmper.

What the research agrees on. The sources differ in kind, from a survey of North American CFOs to an event summary in Belgium, a vendor announcement, and standard setting commentary, and they should not be read as a single body of evidence. They agree on three points. Interest in AI for finance is high. Adoption in the typical function is earlier than the interest suggests. And the control community, through COSO and the PCAOB staff, expects that automated and AI assisted processes will be governed by evidence of operating controls and not by policy statements. The Deloitte close overview and the PwC announcement are the weakest evidence in the table because they come from firms describing their own offerings, and the reader should weigh them accordingly.

8. The ten question self assessment

Rate each statement from 1 to 5 on **demonstrated capability**, using the five point scale of the program: 1 Unprepared (limited understanding, no meaningful application), 2 Developing (conceptual familiarity, relies heavily on others), 3 Competent (can independently apply the capability in conventional situations), 4

Advanced (integrates the capability into strategy and delivers measurable results), and 5 Future-ready (demonstrates enterprise wide impact, anticipates change, and helps others develop). Mark one box in each row. Ratings must rest on actual experience or documented evidence, such as a close calendar you changed or a control you tested, and not on confidence.

Exhibit 14 Module 2 Self Assessment: Rate Each Statement from 1 to 5

No.	Statement	1	2	3	4	5
1	I can describe where my finance function sits on a maturity path and name the evidence for that position.					
2	I can explain how a continuous close differs from a faster monthly close and identify the critical path of my own close.					
3	I can compute days saved and hours saved from a process change and state how much of the hours became a real saving.					
4	I measure forecast accuracy and bias, and I compare every model with a naive benchmark on our own history.					
5	I can explain driver based planning and rolling forecasts and show where each is used in my company.					
6	I can specify the controls that must accompany an automated process, including access, change control, logging, and review thresholds.					
7	I can explain how segregation of duties applies to bots and AI agents and I have tested it.					
8	I can choose among process mining, robotic process automation, rules, and AI agents for a given process and justify the choice.					
9	I can build a wave based business case with payback, a downside case, and gates, and I have presented one to a board or committee.					
10	I report finance performance with measures such as cycle time, touchless rate, accuracy, redeployment, and cost to serve, and I act on them.					
	Total score (maximum 50)					

Rating scale: 1 Unprepared, 2 Developing, 3 Competent, 4 Advanced, 5 Future-ready. Mark one box in each row and add the ten ratings.

Interpreting the total. The bands below are proposed program standards. They are an educational framework and have not been empirically validated against a population of finance executives.

Exhibit 15 Interpretation Bands (Proposed Program Standards)

Score out of 50	Band	What it suggests
10 to 19	Unprepared	Limited experience with finance transformation. Begin with the baseline and the close calendar of one process

Score out of 50	Band	What it suggests
20 to 29	Developing	Conceptual familiarity. Practice on one measured improvement and learn to read a critical path
30 to 39	Competent	Can run a conventional improvement. Strengthen the controls and the downside case
40 to 45	Advanced	Delivers measurable results across processes. Focus on assurance, workforce, and sustaining the gains
46 to 50	Future-ready	Enterprise wide impact and the ability to develop others. Share the method and mentor peers

Source: proposed program standards for the Future-Ready CFO Maturity Index, not empirically validated. Module 2 is not designated a critical module.

Development actions by low score pattern. If statements 1, 2, and 3 score low, the reader lacks a measured baseline. Action one is to measure the close calendar of one process for two cycles. Action two is to draw the dependencies and compute the critical path as in Exhibit 4. Action three is to restate any claimed saving as dollars removed from the budget and hours redeployed. If statements 4 and 5 score low, planning is the weakness. Action one is to compute MAPE and bias for the last twelve months of the main forecast, and to compare it with a naive forecast. Action two is to pick one plan line and tie it to a named driver with an owner. Action three is to hold a bias review with the business units that produce the forecast.

If statements 6 and 7 score low, controls are the weakness. Action one is to list every bot, macro, and agent in use and to name an owner for each. Action two is to add automated identities to the segregation of duties matrix and test the conflicts. Action three is to ask the external auditor what evidence it would need to rely on one automated control. If statements 8, 9, and 10 score low, the weakness lies in selection, economics, and measurement. Action one is to run process mining or a two week sample on one process before choosing a tool. Action two is to rebuild a past business case with a 70 percent downside and gates. Action three is to publish the Exhibit 10 measures for one process each quarter.

9. Currency check and sources

Currency check: the publications in Exhibit 13 were opened and read through the tool described in Section 7, and they are dated between 2021 and 2026. COSO guidance, PCAOB staff views, and vendor offerings in this field change often, and the reader should confirm the current text before relying on any of it. The following statements come from sources that were read: the COSO guidance title, release date of February 23, 2026, six step roadmap, and control building blocks (Deloitte DART, April 3, 2026); the reliance threshold and the pitfalls of paper controls (RSM US, April 7, 2026); the capability view and the advice to consider deterministic automation first (EisnerAmper, May 13, 2026); the staff observations on generative AI risks and uses (PCAOB spotlight, July 2024); the survey statistics of Deloitte CFO Signals Q4 2025 (reported by the Journal of Accountancy, January 13, 2026, and not read in Deloitte's own publication); the KPMG Belgium survey of 157 CFOs; the PwC announcement of May 16, 2025; and the Deloitte close overview with its 2021 copyright. The following come from the author's own knowledge and were not confirmed in a source read: the definitions of continuous close, critical path, process mining, RPA, AI agents, MAPE, bias, WAPE, and the naive forecast; the five stage maturity path, which is a teaching framework and not a standard; the design of the control table in Exhibit 5 beyond the COSO building blocks; and the description of driver based planning, rolling forecasts, controls monitoring, and anomaly detection. COSO itself, NIST, and PCAOB standards were not opened. All companies, people, and figures in the exhibits, including Harbor Ridge Industrial, Inc. and Alder Peak Components, Inc., are fictional, and every calculation was computed in Python. This module is educational. It is not accounting, tax, legal, securities, or professional advice, and the reader should consult qualified professionals in accounting, tax, legal, and other relevant services, and the external auditor, before acting.

Sources: Deloitte DART, COSO Releases Publication on Internal Controls Related to Generative AI (April 3, 2026); RSM US, Is generative AI being used to make decisions you cannot account for? (April 7, 2026); EisnerAmper, COSO AI governance and internal control article (May 13, 2026); PCAOB, Spotlight on generative AI in audits and financial reporting (July 2024); Journal of Accountancy (January 13, 2026); KPMG Belgium, CFOA Next Gen Finance (September 2025); PwC, Agent Powered Performance (May 16, 2025); Deloitte, Next Gen Close (copyright 2021). URLs are in Exhibit 13.