

eFuturesCFO Masterclass Series

AI Workflows for the Modern CFO

PART 3

The Memo

John Campbell's Architectural Framework

Hindol Datta

Chief Content and Product Officer, eFuturesCFO

A Note Before Part 3

Part 2 was the listening. Part 3 is the writing. John Campbell has now spent thirty days conducting structured conversations across the executive team at Helix Cloud Systems. He has heard the patterns. He has identified the workflows that will become the five use cases at the heart of the masterclass. He now has to convert what he has heard into a written architectural framework that Sarah has explicitly requested as a memo, not a slide deck.

What follows is that memo. It is a working document, written for an audience of one and intended to be read carefully. It is the kind of document that serious finance leaders produce in the first sixty days of a new role, often without ever sharing it beyond a small circle, because the act of writing is itself the act of thinking. John is not preparing a deliverable. He is organizing his own mind on the page.

The reader's job in this part is to watch the executive synthesis happen. The discovery conversations of Part 2 were rich but disordered. The memo of Part 3 imposes a structure on what was heard, names the architectural principles that will govern every subsequent decision, justifies the selection of the five use cases against the alternatives, and commits to a sequence. Every section of the memo will be referenced by Parts 4 through 10. The framework that John writes here is the constitution of the masterclass.

The voice shifts from Part 2. Dialogue gives way to executive prose. The memo reads as it would actually read inside a Series B finance function: dense, precise, occasionally first-person, occasionally normative. John is not writing for a general audience. He is writing for Sarah, with the audit committee chair in the back of his mind. The reader is reading over his shoulder.

Read this part slowly. The structure of the memo is the structure of executive judgment under conditions of complexity. The way John moves from observation to thesis to principle to selection to sequencing is itself a transferable practice. The five use cases that emerge will be familiar from Part 2. What is new is why these five, in what order, under what governance, and with what specific commitments. That is what a memo is for.

The act of writing is the act of thinking. A memo is the form in which executive judgment becomes visible to itself.

Hindol Datta

Contents

Memo Header	To, From, Date, Subject
Section 1	The Operating Thesis
Section 2	The Systems Map
Section 3	The Inventory of Pain
Section 4	The AI Readiness Assessment
Section 5	Architectural Principles
Section 6	The Prioritization Framework
Section 7	The Five Use Cases and Why These Five
Section 8	The Sequencing Decision
Section 9	Personal Commitment for the First Hundred Days
Appendix A	Glossary of Internal Terms
Appendix B	Assessment: Twenty Questions
Appendix C	Answer Key with Explanations

What follows is the memo as John writes it. The reader sees the document as the executive sees it: an internal working document on Helix Cloud Systems letterhead, addressed to the CEO, dated thirty-three days into his tenure.

MEMORANDUM

TO: Sarah Chen, Chief Executive Officer
FROM: John Campbell, Chief Financial Officer
DATE: Day 33
SUBJECT: AI in Finance: Architectural Framework and the First Five Workflows
CLASSIFICATION: Internal · Eyes Only

Sarah,

This memo is the deliverable I committed to in our first conversation. It is the architectural framework for the role of artificial intelligence in the Helix finance function and in the broader enterprise where finance has visibility or authority. It draws on twenty-eight days of structured conversations with the executive team and on my own analysis of what those conversations revealed.

I have organized the memo into nine sections. The first six establish the thesis, the system, the pain, the readiness, the principles, and the prioritization. The seventh and eighth name the workflows we will build and the order in which we will build them. The ninth records what I personally commit to in the first hundred days.

I have written this in the voice and at the length I would want it read. There are no slides. There are no bullets. There is no executive summary. You can read the memo in roughly forty minutes. Reading it that way is, I believe, the right investment of time, because what we are about to undertake will shape the next eighteen months of the company.

I look forward to your reactions.

John Campbell

Section 1 · The Operating Thesis

Helix Cloud Systems is, in my reading, a company whose finance function is being asked to see signals that live outside finance. The board's concern about predictability, the recurring pattern of late-quarter forecast surprises, the unexpected churn that Priya catches only forty percent of the time in advance, the persistent infrastructure waste that Tom estimates at six to eight hundred thousand dollars annually, and the close that consumes more of David and Elena's capacity than it should, and each of these is a different surface of the same underlying condition. The company generates signal in many places. Finance sees a fraction of it. The synthesis that finance needs to produce, and that the board increasingly demands, requires reading signal from systems where finance currently has no presence.

This is the thesis I want you to accept before reading anything else in this memo. The finance function does not need to become better at what it currently does. It needs to become structurally different. The work the finance function performs in mid-2026 is largely the work it performed in 2020, with marginal improvements in tooling and process. The demands placed on the function, by contrast, have shifted materially. The board wants forecasts that incorporate leading indicators the current process does not collect. The audit committee wants governance over a technology that did not exist in the function three years ago. The next investor will want operational metrics whose sophistication is far ahead of where we are today. The gap between what the function does and what it must do cannot be closed by working harder. It can only be closed by working differently.

The working differently, in our specific case at our specific stage, is centered on artificial intelligence. I want to be careful in this assertion. I am not claiming that AI is the answer to every operational problem in finance. I am not claiming that AI will replace any of the people you have inherited. I am claiming, narrowly and specifically, that AI is the technology that compresses analytical work enough to make multi-source signal aggregation practical at our scale and stage. We do not have the bandwidth to build a data science function. We do not need to. We need a small number of well-architected workflows, governed properly, that bring AI to the signals that already exist in our systems.

What the thesis implies

Four implications follow.

First, finance becomes the place where the fragments come together. The customer health signals that live in Priya's function, the rep behavior signals that live in Michael's function, the product usage signals that live in Raj's systems, the macro and operational signals that live in Karen's and Tom's work, these flow into the finance function because finance is the natural locus of synthesis. The board is asking finance for synthesis. The synthesis cannot happen until the fragments are gathered. The gathering is what we will design.

Second, AI in finance is a governance question before it is a technology question. The technology will work, by which I mean the models will perform their tasks competently if we configure them well. The governance will not work unless we build it deliberately. The CFO's office is the right home for the governance because the consequences of AI failure in finance land in the CFO's office regardless of where the failure originates. Naomi made this argument to me directly in our conversation and I agree with it.

Third, the right architecture is workflows rather than agents. For the next eighteen months, every AI deployment we sanction in finance will be a workflow with a defined sequence and a named human owner. Autonomous agents will be considered, but only with a specific case made for each one. The default posture is human-in-the-loop, deterministic flow control, comprehensive audit trail. This is conservative. It is also correct for our stage and for the audit committee's expectations.

Fourth, the five use cases we will build are not a list. They are a system. Each of them addresses a fragment of the same underlying pattern. The forecasting workflow draws customer signals into the revenue view. The board reporting workflow compresses the reporting cycle. The pipeline intelligence workflow brings rep behavior patterns into the forecast. The finance operations workflow reclaims time from the close. The infrastructure and vendor intelligence workflow brings continuous monitoring to spend categories where episodic monitoring has been allowing waste to accumulate. Built together, the five workflows produce the synthesis the board has been asking for. Built individually, they produce efficiency gains but not the structural change in what the finance function is.

What the thesis does not claim

It is worth saying clearly what the thesis does not claim.

It does not claim that AI will improve our forecast accuracy simply by being deployed. The forecast will improve only if the workflow draws on signals that the current process does not collect. The model is not magic. The signal sources are what matter.

It does not claim that the close can be eliminated. The close will be compressed by automating mechanical work, but the judgment, the controls, and the financial reporting itself will remain the responsibility of human accountants. The model reduces the routine. The accountants retain the substance.

It does not claim that the eighteen-month horizon is sufficient for transformation. We will be markedly better in eighteen months. We will not be fully where we should be. The next CFO of Helix, whether that is me in a different role or a successor, will inherit a function that is mid-journey. That is acceptable. The objective is to start the journey deliberately, not to complete it on a calendar.

The thesis in one sentence

Finance at Helix is a multi-source signal aggregation problem that the current architecture is not solving, and AI is the technology that makes the aggregation practical at our scale and stage, governed by the CFO's office under conservative workflow patterns.

Section 2 · The Systems Map

Before proposing any architecture, I want to describe the current-state system as I now understand it. The discovery conversations gave me visibility into where data lives, how it moves, where it gets transformed, and where the handoffs occur. The picture is reasonably clean for a company of our stage, with notable gaps that the new architecture must address.

Systems of record

Five enterprise systems hold the data that finance depends on most directly.

Salesforce is our customer relationship management system. It holds the master record of opportunities, contacts, and accounts. Michael's team operates with reasonable discipline at the deal level, less discipline at the contact level. The opportunity data is reliable for current state. The opportunity history is reliable for the past two years but less reliable before that because of stage definitions that have evolved.

NetSuite is our enterprise resource planning system. It holds the general ledger, accounts receivable, accounts payable, and most of the operational financial data. David and his team operate it with discipline and the data is clean. The audit has been clean in each of the past two years.

Rippling is our human resources information system. It holds the headcount roster, the org structure, compensation data, and the hiring pipeline. Karen has stabilized this system in the past eight months. The data is reasonably reliable.

Gusto holds benefits administration. ADP holds payroll. Pulley holds the cap table. Each of these is a smaller system in our integration footprint but each holds data that the finance function eventually consumes.

The product application itself, running on AWS in us-west-2 and eu-west-1, generates substantial event-level data about how customers use Helix. This data flows into Snowflake but is voluminous and inconsistent across product versions because we have changed our event schema twice since instrumentation began.

The data warehouse

Snowflake is the current convergence point. Fivetran connectors pull from Salesforce and NetSuite on a regular schedule. A custom pipeline pulls product event data from the application tier. Rippling data is pulled less reliably. The smaller systems are not currently in the warehouse.

The warehouse is real and is being used. Looker sits on top of it and serves the dashboards. Elena has authored most of the analytical content in Looker, despite the warehouse being officially owned by a data engineer in Raj's organization. The practical owner of analytical content is therefore the senior FP&A;

manager, which is a structure I will need to address.

The warehouse data quality is variable by source. Salesforce data at the deal level is high quality. Contact-level Salesforce data is low quality. NetSuite data is clean. Product event data is voluminous but inconsistent. Headcount data from Rippling is reasonable. The smaller systems are not in the warehouse and their data must be retrieved manually or through one-off scripts.

The current AI footprint

Helix currently runs two sanctioned AI deployments and an unknown number of unsanctioned ones.

The sanctioned deployments are GitHub Copilot on the engineering team, with enterprise data handling commitments and proper governance, and the Cadenza pilot in revenue operations, currently on Michael's personal expense card at approximately forty-two hundred dollars per month. The Cadenza deployment has a data processing agreement but the agreement was not reviewed by counsel before signing.

The unsanctioned uses include marketing's team subscription to ChatGPT, customer success's use of Claude for customer-facing messages, Elena's use of ChatGPT for variance commentary on a personal subscription, Raj's and several engineers' use of Claude and ChatGPT for various engineering tasks, and probable but unconfirmed uses in finance and elsewhere. There is no central inventory. There is no policy framework. There is no audit trail.

The single most important finding

The current AI footprint at Helix is not the AI footprint we think it is. Sanctioned deployment is small. Unsanctioned deployment is broad. The governance gap is the largest single risk we currently carry in the finance and operational perimeter, and it must be closed before any new deployment begins.

The decision flows

Three decision flows matter most for the finance function.

The forecast flow is monthly. Elena builds the forecast over three days from pipeline data, headcount plans, and operational cost projections. The forecast is reviewed by David and then by me. The forecast is reconciled against the board plan and distributed in the close package. The forecast is systematically optimistic on revenue in the final weeks of each quarter by three to four percent on average, correlating with rep staging behavior that has not been disciplined.

The close flow is six business days. Day one is cutoff and bank reconciliations. Day two is subledger close and preliminary numbers. Day three is accruals and allocations, where most of the time goes. Day four is management review. Day five is final statements and operational reporting. Day six is package distribution. The close is competent but operates close to the capacity ceiling of David and Elena, who are at risk of

burnout.

The board package flow is approximately one week, concentrated in the two days before the board meeting. The package is currently constructed separately from the close rather than as a derivative of it. Inputs come from multiple functions and are integrated by David. The package contains roughly twenty-five to thirty slides and an exhibit pack.

The governance and security perimeter

We have a SOC 2 Type 2 report that is current. We have SSO, MFA, and least-privilege production access. We have audit logging on production systems. We do not have an AI tool inventory. We do not have a data flow map for AI systems. We do not have an approval process for new AI tools. We do not have a policy framework for employees on what AI use is permitted. We have informal practice, which is not a control.

Naomi has flagged that the regulatory environment for AI has shifted enough in the past twelve months that even sophisticated observers need to update their mental model. The EU AI Act is in force. Several US states are moving faster than the federal government. Our enterprise customers, particularly in regulated industries, are beginning to include AI clauses in their procurement agreements. We can answer few of these cleanly today.

This is the system as it exists. The architecture I will propose builds on what is good (the warehouse, the close discipline, the SOC 2 posture, the capacity model Karen has built) and addresses what is missing (the AI governance, the multi-source synthesis, the leading-indicator forecasting, the continuous monitoring of operational spend).

Section 3 · The Inventory of Pain

Before designing the architecture, I want to be specific about the operational pain that the discovery conversations surfaced. Each item below is a workflow that is currently broken, slow, or unreliable. Each one has been identified by a specific function leader and is documented in my notes. I list them in approximate order of business impact.

Forecast unreliability at quarter-end

The forecast misses by three to four percent in the final four weeks of each quarter, systematically over-forecasting rather than under-forecasting. Two-quarter and three-quarter patterns of these misses have produced the late-quarter surprises that the board cites as its primary concern. The cause is twofold: rep staging behavior that advances deals to late stages based on relationship confidence rather than on objective gating criteria, and a forecast architecture that reads only pipeline data rather than incorporating leading indicators from customer success, product usage, and rep behavior history. Impact: direct board predictability concern, direct risk to Series C terms in twelve to fourteen months. Owner: Elena Vargas builds the forecast, Michael owns the pipeline behavior.

Churn surprise rate of sixty percent

Approximately sixty percent of accounts that churned in the past twelve months had been rated green or yellow in their last health review. The health scoring is produced by CSMs from subjective assessment rather than from aggregated signals across product usage, support ticket patterns, engagement cadence, and billing history. The signals exist. The aggregation does not. Impact: indirect revenue loss through unanticipated churn, direct distortion of NDR and forecast assumptions, erosion of customer success function credibility. Owner: Priya Sundaram.

Six-day close consuming finance capacity

The close runs six business days against a peer benchmark of four to five days for a company of our size. Day three consumes a full day and a half on accruals and allocations that are mechanically intensive. The marketing accrual alone takes half a day. Approximately two hundred to four hundred ambiguous GL transactions per month consume an additional full day of David's time on categorization review. The function is operating close to capacity ceiling and is at risk of burnout. Impact: direct constraint on the finance function's ability to take on additional work, indirect pressure on Diana's requested five-day close, retention risk on the two strongest individuals in the function. Owner: David Kim and Elena Vargas.

Board package decoupled from the close

The board package is constructed as a separate exercise from the close, consuming approximately one week of effort concentrated in the two days before the board meeting. Inputs come from multiple functions and are integrated by David. The variance commentary is drafted manually. The board commentary is drafted manually. Each board cycle is essentially built from scratch rather than as a derivative of the underlying data and the prior period's template. Impact: direct consumption of David's time during the highest-leverage week of each quarter, indirect quality variance across board cycles, missed opportunity to deliver the kind of visibility Sarah has asked for. Owner: David Kim with contributions from each function head.

AWS infrastructure waste at twenty to twenty-five percent

Tom estimates that twenty to twenty-five percent of our AWS spend is waste rather than legitimate consumption, representing sixty to eighty thousand dollars per month or seven hundred twenty thousand to nine hundred sixty thousand annually. The waste accumulates because there is no continuous monitoring of utilization with accountability at the team level. Engineering cleanups happen episodically and the pattern of growth resumes within two months of each cleanup. Impact: direct cash burn, indirect distortion of unit economics. Owner: shared between Tom Reeves and Raj Patel.

Software license utilization gaps

Approximately twenty-nine percent of our SaaS license portfolio is paid for and unused, representing approximately one hundred eighty thousand dollars annually. The number is incomplete because Tom tracks utilization manually through quarterly surveys with a response rate of approximately eighty percent and variable accuracy in the responses. The true waste is likely higher. Impact: direct cash burn, cumulative across many small vendor relationships. Owner: Tom Reeves.

Pipeline behavior patterns invisible to finance

Reps systematically advance deals into late stages based on confidence rather than on gated criteria. Stage 5 slip rates have not been measured over the eight-quarter window I requested. The pattern is visible from the FP&A; side but has not been addressed at the sales operations level. Impact: direct contribution to forecast unreliability, indirect erosion of executive trust in the pipeline. Owner: Michael O'Brien with sales operations support from Lisa Mahoney.

Shadow AI usage across the enterprise

An unknown number of AI deployments exist outside the sanctioned perimeter. Marketing has a team ChatGPT subscription. Customer success has been using Claude for customer messages. Elena uses ChatGPT for variance commentary on a personal account. Multiple engineers have personal subscriptions. A customer success subscription may or may not still be active. The Cadenza pilot is on a CRO's personal expense card. No central inventory exists. No policy framework exists. No data flow map exists. Impact: governance and regulatory risk, potential customer contractual exposure, audit committee concern. Owner: governance gap, currently unowned, will be assumed by the CFO's office.

NDR reporting at one hundred eighteen percent versus reality of one hundred thirteen

Our reported net revenue retention is one hundred eighteen percent. Approximately thirty percent of that growth is contractual escalators and price increases at renewal rather than true expansion bookings. The honest expansion-only NDR is closer to one hundred thirteen percent. The reporting has not been corrected because the prior interim CFO did not adjust it. Impact: direct misrepresentation risk in next investor diligence, indirect erosion of metric credibility once discovered. Owner: David Kim and the new CFO.

AI resume screening paused on regulatory grounds

Karen paused the talent acquisition team's AI resume screening tool because of regulatory exposure under New York City, California, and EU AI Act provisions covering automated employment decision tools. The pause is correct. The absence of a clear classification framework for which AI uses fall into high-risk regulatory categories is a broader gap that extends beyond hiring. Impact: present-tense compliance exposure on any future deployment in regulated categories. Owner: Naomi Bridges with the CFO's office.

Pattern

The pattern across the ten items is consistent. In each case, a functional leader has visibility into the problem, has partially diagnosed it, and lacks either the tools or the authority to address it structurally. The finance function is the natural place to consolidate the structural response because the finance function is where consequences land and because the audit committee has signaled it wants the CFO to own the AI governance perimeter.

Not all ten items will be addressed in the first eighteen months. The five use cases I propose in Section 7 address the items with the largest aggregate impact and the most reachable architectural solutions. Several remaining items will be addressed through governance, policy, and process changes rather than through new AI workflows. The NDR restatement, the AI resume screening situation, and the pipeline behavior discipline are examples of issues that require executive judgment and process change more than they require new

technology.

Section 4 · The AI Readiness Assessment

Most companies overestimate their AI readiness. The discovery process suggests Helix is closer to ready than the average Series B company but is not as ready as the executive enthusiasm for AI would imply. This section assesses readiness honestly across five dimensions. The purpose is not to discourage deployment. The purpose is to identify which foundational gaps must close before serious workflows can operate on the foundation.

Data readiness

Data readiness is moderate. The warehouse exists and is meaningfully used. Salesforce, NetSuite, and Rippling data flow into it on a reliable schedule. Product event data flows but is inconsistent across version boundaries. Smaller systems are not in the warehouse. Data quality is highest in NetSuite, next in Salesforce at the deal level, weaker at the contact level and the product event level.

The implication for sequencing is that the workflows that rely on NetSuite and on deal-level Salesforce data can be built sooner. The workflows that rely on product event data, on contact-level Salesforce data, or on data from systems not currently in the warehouse will require data preparation work that adds weeks or months to the timeline.

Governance readiness

Governance readiness is poor. We have informal practice rather than a written policy. We have no AI tool inventory. We have no data flow map for AI systems. We have no approval process for new AI deployments. We have no defined risk tiering for AI use cases. We have shadow usage across the enterprise that we cannot currently inventory.

The implication is that governance work must precede deployment work. The first sixty days of substantive activity after this memo is approved will be spent on the governance framework rather than on the workflows themselves. This is the correct sequence, and it is the ordering Diana endorsed in our conversation.

Engineering capacity

Engineering capacity for the foundation work is available but constrained. Raj has indicated that he can free one senior engineer for the three foundational components he proposed: the MCP server in front of the warehouse, the workflow registry, and the audit log. He has committed to a usable version of all three within three months at one engineer dedicated, or an enterprise-grade version within six months. We will need both.

Beyond the foundation work, individual workflows will require engineering time to implement. My estimate, based on comparable builds in my prior CFO roles, is that each workflow consumes approximately six to ten engineer-weeks from design through production deployment. Across five workflows that is thirty to fifty engineer-weeks, which represents a meaningful share of Raj's annual engineering capacity. We will need to be deliberate about sequencing and about which workflows justify the engineering investment.

Finance team capacity

Finance team capacity is the most constrained variable in the assessment. David and Elena are operating near their capacity ceiling. Adding workflow design, vendor evaluation, and governance development to their plates is not feasible without either adding headcount or reducing the work they currently perform.

My intention is to address this through three actions in parallel. First, hire one additional FP&A; analyst within sixty days, giving Elena the bandwidth to lead the forecasting workflow design without dropping current responsibilities. Second, deploy the finance operations copilot early so that the close work compresses for David. Third, define Elena's role explicitly as the FP&A; lead with AI workflow design authority, recognizing the work she has been doing informally. The combination of headcount, tooling, and structural recognition addresses the capacity risk without expanding the headcount budget materially.

Vendor and contractual readiness

We are not currently in a position to defend our AI deployments to a sophisticated enterprise customer or to a Series C investor's technical diligence. We do not have a subprocessor list for our AI use. We do not have data processing agreements with all our AI vendors. We do not have audit trail evidence for our AI-driven outputs.

Closing these gaps requires legal, security, and procurement work that is largely in flight or near-term. Naomi has committed to drafting the policy framework. Wei has committed to designing the technical controls. Tom will own the vendor inventory and the renegotiation of the Cadenza relationship onto a master agreement. The work is concrete and the owners are named. The timeline is approximately ninety days for the core artifacts.

Overall readiness

Overall, Helix is in approximately the upper third of Series B companies on AI readiness, primarily because the engineering foundation is more mature than typical. We are in approximately the lower third on governance readiness, because the governance work has not begun in earnest. The arithmetic of readiness is that the lowest score dominates. A company with strong data and weak governance is not ready for serious AI deployment in finance, because the governance gap will produce the failure that ends the program.

The readiness principle

Readiness is determined by the weakest dimension, not by the average. Until governance readiness rises to meet data and engineering readiness, we are not ready to deploy production AI workflows in finance.

Section 5 · Architectural Principles

The next seven principles are the rules I propose to apply to every AI deployment in the Helix finance function. They are short and declarative because principles must be reusable. I have drafted them in a form that I expect to publish, with your approval, as the operating constitution for AI in finance.

Principle One: Workflows before agents

Every AI deployment in finance will be designed as a workflow with deterministic flow control and a defined sequence of steps, unless a specific case is made for an autonomous agent and that case is approved by the CFO and the audit committee chair. Workflows are auditable, predictable, and easier to govern. Agents are powerful but harder to control. For a finance function in 2026, workflows are the default and agents are exceptional.

Principle Two: Human-in-the-loop for all financial outputs

No AI-generated output that affects the financial statements, the board package, customer-facing financial communications, or any external reporting will be released without a named human reviewer. The reviewer is responsible for the output, not the model. The audit trail will record both the model's output and the human review.

Principle Three: No production AI without an audit trail

Every AI workflow deployed in finance will produce an immutable audit trail capturing the input, the model version, the prompt or workflow definition, the output, the human reviewer if any, and the time stamps. The audit trail will be retained for the period required by our audit and regulatory obligations. No exceptions.

Principle Four: No sensitive data leaving approved environments

No customer personally identifiable information, no employee compensation data, no acquisition or transaction-sensitive information, and no information classified as confidential will leave the approved AI environment without a specific authorization. The approved environments are defined by Wei and Naomi and will be updated as our vendor list evolves.

Principle Five: No vendor lock-in without a substitution path

Every AI vendor relationship and every AI workflow we deploy will be designed so that the underlying model can be substituted if the vendor relationship changes. Where possible, the Model Context Protocol will be used for integrations rather than custom connectors. Where MCP is not yet available, the substitution path will be documented before the deployment proceeds.

Principle Six: Every workflow has a named human owner

Every AI workflow in production will have a single named human owner accountable for its operation, its output quality, its compliance with the policy framework, and its periodic review. Ownership cannot be a committee. Ownership must be a person.

Principle Seven: Governance precedes deployment

No new AI workflow will be deployed in finance until it has been classified under the risk tier framework, reviewed by the governance working group, approved by the CFO, and documented in the workflow registry. No exceptions, including for workflows that appear low-risk. The discipline of governance precedence applies uniformly because exceptions undermine the framework.

The principles in summary

Workflows over agents. Human-in-the-loop for financial outputs. Audit trail for everything. No sensitive data leakage. No vendor lock-in. Named owners. Governance before deployment. Seven principles, no exceptions.

These seven principles are deliberately conservative. They will, in some cases, slow our deployment relative to what is technically feasible. That cost is acceptable. The cost of a single high-profile AI failure in finance is greater than the cost of moving slowly. Diana has made clear that the audit committee will be patient with the pace and impatient with shortcuts. The principles are designed to honor that posture.

Section 6 · The Prioritization Framework

A finance organization can deploy AI in many places. The discipline of selection is one of the highest-leverage executive activities, because deploying the wrong workflow first sets the program back by months and erodes the organization's appetite for subsequent deployments. I propose five criteria for evaluating AI use case candidates and three thresholds for sequencing among the qualifying use cases.

The five criteria

Each criterion is a question I ask of every candidate workflow. A workflow that fails on any one of them is either not yet ready for deployment or not a good candidate at our stage.

1. Business impact

Does this workflow address a problem with measurable financial impact, executive visibility, or board-level concern? Workflows that produce efficiency gains in invisible corners of the operation are not the right place to begin. The first workflows must produce value that registers at the executive level, both to justify the investment and to sustain the organizational appetite for the next deployment.

2. Executive sponsorship

Is there a named functional leader who has explicitly requested this workflow, will own its deployment, and will consume its output? Workflows imposed on functions whose leaders did not ask for them tend to fail. Workflows owned by leaders who requested them tend to succeed.

3. Data readiness

Are the data sources required by this workflow already in a usable state, or will substantial data preparation work precede deployment? Workflows that depend on data that does not yet exist in usable form must wait until the data preparation work is complete, regardless of other merits.

4. Governance risk

Is this workflow sufficiently contained that it can be deployed under the governance framework without requiring expansion of the framework? Workflows that push against the edges of the governance framework should not be the first deployments. The first deployments should be inside the framework, not at its boundaries.

5. Reversibility

If this workflow produces a wrong output, can the wrong output be detected and reversed before consequential harm occurs? Workflows whose errors cannot be easily reversed should be deployed later, after the team has accumulated experience with lower-risk workflows.

The three sequencing thresholds

Among the workflows that pass the five criteria, three thresholds determine sequencing.

First threshold: foundation dependency

Workflows that build on the foundation must wait until the foundation is in place. The MCP server in front of the warehouse, the workflow registry, and the audit log must be operational before workflows that depend on them can be safely deployed. This is the constraint that sets the earliest possible deployment date.

Second threshold: governance maturity

Workflows whose risk tier exceeds the maturity of the governance framework must wait. The earliest workflows should be the ones whose risk profile is well-matched to the governance maturity at the time of deployment. As the governance framework strengthens, higher-risk workflows become eligible.

Third threshold: executive bandwidth

The organization can absorb a limited number of new deployments simultaneously. My estimate is that we can absorb one major workflow deployment per quarter, plus continuing operation of previously deployed workflows. Attempting two simultaneous major deployments overloads the engineering team, the finance team, and the governance review process.

The framework applied

Each of the five use cases I propose in Section 7 has been evaluated against the five criteria and against the three sequencing thresholds. The sequencing decision in Section 8 is the output of that evaluation. The framework itself is reusable, so when new candidate workflows emerge in subsequent quarters, the same evaluation can be applied.

Section 7 · The Five Use Cases and Why These Five

I propose to deploy five AI workflows in the Helix finance function over the next eighteen months. The five were not chosen as a list. They emerged from the discovery conversations and were filtered through the prioritization framework in Section 6. Each addresses a specific pain documented in Section 3. Each is appropriately scoped for our stage. Each has named executive sponsorship. Each can be governed under the framework I will propose to the audit committee.

In this section I describe each use case briefly: the problem it addresses, the architecture in summary, the success criteria, the principal governance considerations, and the conversation it traces back to. The detailed design of each will be the subject of one of the masterclass parts that follow, after the governance framework is established in Part 4. This section is the executive justification, not the technical specification.

Use Case One: AI-Augmented Forecasting Engine

Problem. The current forecast is built bottom-up from pipeline data and headcount plans. It misses by three to four percent in the final four weeks of each quarter, systematically optimistic, contributing directly to the late-quarter surprises the board has cited. The forecast does not currently incorporate leading indicators from customer success, product usage, rep behavior history, or macro conditions.

Architecture in summary. A workflow that aggregates signals across the existing systems (Salesforce, the warehouse, customer success notes, hiring pace, macro indicators), applies AI-assisted analysis to identify which signals are most predictive in the current period, and produces a rolling forecast with confidence intervals, driver analysis, and variance commentary. The forecast remains human-owned. The AI augments rather than replaces the FP&A; judgment. Elena leads the design.

Success criteria. Forecast accuracy on the four-week-out revenue view improves from current five to eight percent error to three percent or better. Late-quarter surprises reduce in frequency and magnitude. The forecast incorporates leading indicators that the current process does not see.

Governance considerations. The forecast itself is an output consumed by the board and the next investor. The risk tier is high. Human-in-the-loop review is non-negotiable. The audit trail must capture which signals fed the forecast, how they were weighted, and what human judgments were applied. Reversibility is good because forecasts are updated monthly and errors in one cycle are observable in the next.

Traces to. Sarah's framing of predictability as the board's primary concern. Michael's acknowledgment of rep staging patterns. Elena's articulation of the leading-indicator gap. Diana's deliverable of three-percent accuracy on the four-week-out view.

Use Case Two: AI Board Reporting and Narrative System

Problem. The board package is currently constructed as a separate exercise from the close, consuming approximately one week of David's time concentrated in the two days before the board meeting. Inputs come from multiple functions and are integrated manually. Variance commentary and board commentary are drafted from scratch each cycle.

Architecture in summary. A workflow that ingests the close data as soon as it is available, draws on prior period commentary as a stylistic reference, and produces draft board commentary, KPI summaries, variance explanations, and risk callouts. The draft is reviewed and edited by the finance team. The board package becomes a derivative of the close rather than a separate construction. David leads the design.

Success criteria. Board package construction time reduces from one week to two days. Board commentary quality improves through consistency across cycles. Variance commentary becomes more analytical and less descriptive. Elena recaptures the time she currently spends on draft preparation.

Governance considerations. The board package is an external-facing artifact in the sense that it goes to the audit committee. Every AI-drafted section must be reviewed by the CFO before distribution. The audit trail captures the source documents the AI drew from and the human edits applied. Risk tier is moderate to high because errors in the board package have direct executive consequences.

Traces to. David's description of the board package workflow. Elena's existing informal use of ChatGPT for variance commentary. Sarah's request for a function that sees around corners. Diana's expectation of institutional quality reporting.

Use Case Three: Pipeline-to-Revenue Intelligence System

Problem. Reps systematically advance deals into late stages based on relationship confidence rather than on objective gating criteria. Stage 5 slip rates have not been measured over the relevant historical window. The pattern is visible from the FP&A; side but has not been addressed at the sales operations level. The pattern contributes directly to forecast unreliability.

Architecture in summary. A workflow that analyzes pipeline data continuously, scoring each opportunity's probability of close in the forecasted quarter based on stage history, engagement patterns, age in stage, comparable historical deals, and rep-specific historical reliability. The output is a probability-weighted forecast, slippage alerts on specific deals, and pattern alerts on specific reps. The output is consumed by Michael and the sales operations function, with the finance function as a stakeholder. Michael leads the design with Lisa Mahoney.

Success criteria. The probability-weighted forecast outperforms the rep-stage-based forecast by a measurable margin. Late-quarter slippage is detected earlier in the quarter. Rep-specific patterns of optimism or pessimism are visible and addressable.

Governance considerations. The workflow analyzes sales behavior, which raises sensitivity around how the analysis is used in performance management. The use of the output for forecast intelligence is unambiguous and supported. The use of the output for individual rep evaluation must be governed carefully. Risk tier is moderate. Reversibility is good.

Traces to. Michael's acknowledgment of staging behavior. Elena's identification of the systematic over-forecasting pattern. Sarah's framing of predictability.

Use Case Four: AI Finance Operations Copilot

Problem. The close consumes six business days against a peer benchmark of four to five days. Day three on accruals and allocations consumes most of the total time. Two hundred to four hundred ambiguous GL transactions per month consume an additional day of David's time on categorization review. The finance team is operating near capacity ceiling. Routine analysis is being performed by senior people who should be doing higher-leverage work.

Architecture in summary. A workflow that classifies transactions automatically with confidence scoring, handles routine accrual calculations with auditable derivation, flags exceptions for human review, and provides a copilot interface for finance team members to ask questions of the general ledger and the close data. The copilot is bounded to finance data and operates with full audit logging. David leads the design with Elena.

Success criteria. Close compresses from six days to five. David's monthly review time on ambiguous transactions reduces from one day to two to three hours. The copilot becomes the daily working tool for finance team analytical questions, replacing manual queries and spreadsheet lookups.

Governance considerations. The workflow touches the general ledger and the close. Every automated classification must be reviewable and reversible. Audit trail requirements are strict. Risk tier is high because errors flow into the financial statements. Human-in-the-loop is non-negotiable. Auditor coordination will be required before deployment.

Traces to. David's description of close pain. Elena's articulation of the mechanical work consuming her time. Karen's flag about finance team burnout.

Use Case Five: AI Supply Chain and Infrastructure Intelligence System

Problem. AWS infrastructure waste runs at twenty to twenty-five percent of spend, representing seven hundred twenty to nine hundred sixty thousand dollars annually. Software license waste runs at twenty-nine percent of the license portfolio, representing approximately one hundred eighty thousand dollars annually. Vendor contract terms are reviewed manually at renewal, with limited bandwidth for the long tail of smaller contracts.

Architecture in summary. A workflow that monitors AWS utilization continuously, flagging waste candidates and producing weekly summaries for engineering leadership. A second component aggregates software license utilization from API-accessible sources and prompts for confirmation where API access is not available. A third component analyzes vendor contracts at renewal to flag risky terms, auto-renewal traps, and pricing escalators. Tom leads the design with engineering and procurement support.

Success criteria. AWS waste reduces by at least half within twelve months of deployment, representing three to four hundred thousand annualized. License waste reduces by a similar proportion. Vendor renewals consistently catch risky terms before contract execution.

Governance considerations. The workflow operates on operational and procurement data rather than financial reporting data, so the audit risk is lower. The risk tier is moderate. Reversibility is high because operational spend errors are observable in the next billing cycle. Vendor data confidentiality is the principal sensitivity.

Traces to. Tom's articulation of the AWS waste pattern, the license utilization gap, and the vendor contract review backlog.

Why these five and not others

Several adjacent use cases were considered and deferred. It is worth naming them and explaining why.

A customer health scoring workflow was considered. It is important and it traces directly to Priya's articulation of the sixty-percent churn surprise rate. I have deferred it to a second wave because the architecture overlaps substantially with the forecasting workflow, the data preparation requirements are heavier, and the customer success function does not currently have the bandwidth to lead the design. The capability will be partially captured as a side effect of the forecasting workflow, which will consume customer health signals as an input.

A capacity model dynamic update workflow was considered. It traces to Karen's articulation of the static capacity model. I have deferred it because the existing model is serviceable, Karen has the bandwidth to update it manually on a quarterly cadence, and the marginal value of dynamic updating is lower than the marginal value of the five use cases I have proposed.

An AI-assisted recruiting or talent acquisition workflow was considered. It is the use case that the regulatory environment has rendered most exposed in the near term. Karen has paused the existing tool. I will not deploy anything in this category until the regulatory landscape stabilizes and the governance framework can accommodate high-risk employment use cases.

A customer-facing AI capability inside the Helix product itself was considered. It is the most exciting category on this list. It is also the one that requires product engineering ownership rather than finance ownership, and its development should be sequenced relative to product strategy rather than relative to this memo. I will support Sarah and Raj in this decision but will not propose it from the finance side.

The five in summary

Forecasting. Board reporting. Pipeline intelligence. Finance operations copilot. Infrastructure and vendor intelligence. Each addresses a specific documented pain. Each has named ownership. Each is appropriately scoped. Together, they produce the multi-source synthesis the board has been asking for.

Section 8 · The Sequencing Decision

The order in which we deploy the five workflows is itself a consequential decision. Each deployment teaches the organization something that affects the next. The wrong order can undermine the program even if the individual workflows are well-built. I propose the following sequence, organized by quarter against a baseline of Q1 of the next fiscal year as the starting point. The current calendar context puts us approximately five months from Q1, which is enough time for the foundation work.

Quarter Minus One (current quarter through next): foundation

No workflows are deployed in this period. The activity is governance framework drafting, technical foundation construction (the three components Raj proposed), AI vendor inventory and cleanup of the Cadenza relationship, finance team capacity additions (one new FP&A; analyst), policy publication, and audit committee briefing on the framework. This is the work that makes the deployments possible. It is not glamorous. It is essential.

Quarter One: Finance Operations Copilot (Use Case Four)

The first workflow is the finance operations copilot. The reasoning is that it produces immediate, visible value to the finance team itself, gives David and Elena their first experience operating an AI workflow under governance, reduces the close burden so that capacity becomes available for subsequent workflow design work, and operates on data that is already clean (NetSuite). The risk profile is moderate to high because the workflow touches the general ledger, but the human-in-the-loop discipline is most natural in this use case because finance team members are already accustomed to reviewing transaction classifications. Deploying first into the function that owns the governance framework is the right choice operationally and politically.

Quarter Two: AI Board Reporting (Use Case Two)

The second workflow is board reporting. By Q2, the close has been compressed and the finance team has bandwidth. The workflow builds on the close data that the copilot has already standardized. Elena has experience with the informal version (her variance commentary use of ChatGPT) and can lead the design quickly. The first board meeting in which the new package is used is Q3, which gives us a full quarter of iteration before the package is consumed by Diana and the rest of the audit committee.

Quarter Three: Pipeline Intelligence (Use Case Three)

The third workflow is pipeline intelligence. By Q3, we have data history and we have the foundation in place. The workflow runs in parallel with the existing rep-stage-based forecast for at least one quarter before influencing the official forecast. Michael owns the deployment with Lisa Mahoney; the finance function consumes the output. The parallel run reduces the deployment risk and lets us calibrate the model against actual outcomes.

Quarter Four: Forecasting Engine (Use Case One)

The fourth workflow is the forecasting engine, the most consequential of the five and the one that most directly addresses the board's predictability concern. It is sequenced fourth because it depends on the pipeline intelligence workflow as one of its inputs, because it requires the finance team to have experience operating multiple workflows under governance, and because the failure cost is highest. Deploying the highest-stakes workflow after the team has matured on three lower-stakes workflows is the right risk posture. By Q4 we are roughly eighteen months into my tenure and three months from the Series C window.

Year Two Quarter One: Infrastructure and Vendor Intelligence (Use Case Five)

The fifth workflow is infrastructure and vendor intelligence, sequenced last not because it is unimportant but because it is the most independent of the other four and can be deployed at any time without affecting them. Tom owns it. The value it produces is largely cash savings, which compound regardless of when in the eighteen-month window they begin. Sequencing it last allows the finance function to absorb the first four workflows fully before adding a fifth that operates in a different functional domain.

Alternatives considered

I considered three alternative sequencings.

The first alternative would have led with the forecasting engine because of its direct alignment with the board's predictability concern. I rejected this because deploying the highest-stakes workflow first, before the organization has experience operating any workflow under governance, is the configuration most likely to produce a high-profile failure that ends the program.

The second alternative would have led with infrastructure and vendor intelligence because it is the most independent and the lowest risk. I rejected this because deploying first into Tom's function rather than into finance misallocates the political capital and the executive attention. The first workflow should produce visible value to the finance function itself, both to justify the investment internally and to give the finance team firsthand experience with the governance framework they are owning.

The third alternative would have run two workflows in parallel from Q1 (the copilot and infrastructure). I rejected this because the engineering team, the finance team, and the governance review process can absorb one major deployment per quarter, not two. Attempting parallel deployment in the first quarter introduces unnecessary risk for a marginal acceleration of the program.

The sequence in summary

Q-1: foundation. Q1: finance operations copilot. Q2: board reporting. Q3: pipeline intelligence. Q4: forecasting engine. Year 2 Q1: infrastructure and vendor intelligence. Eighteen months from today to a fully deployed five-workflow portfolio under governance.

Section 9 · Personal Commitment for the First Hundred Days

A framework that lacks personal commitment from the executive who proposed it is a framework that will not be executed. In this final section, I record what I personally commit to doing in the first hundred days, against which I expect to be held accountable by you and by Diana.

What I will do personally

In the first thirty days following acceptance of this memo, I will draft the AI governance framework that will become the operating constitution for AI in finance and a reference point for AI use elsewhere in the company. The framework will be reviewed by Naomi and Wei before it reaches you. The framework will be presented to the audit committee at the next meeting. I will personally write the framework rather than delegate it.

In the first sixty days, I will publish the AI usage policy for the company. The policy will be drafted by Naomi, reviewed by me and Wei, approved by you, and distributed by Karen to all employees. The policy will include the sanctioned tool list, the prohibited uses, the data handling requirements, and the exception request process.

In the first sixty days, I will complete the AI vendor inventory and bring the Cadenza relationship onto a corporate master agreement. The inventory will be maintained by Tom going forward as part of his ongoing vendor management responsibility.

In the first ninety days, I will hire one additional FP&A; analyst, recognize Elena formally as the FP&A; lead with workflow design authority, and ensure that David has bandwidth to participate in the workflow design process.

In the first hundred days, I will present to you and to the audit committee a deployment plan for the first workflow (the finance operations copilot) with engineering commitments, success criteria, governance approval, and a specific deployment date in Q1 of next fiscal year.

What I will not do

It is equally important to record what I will not do.

I will not approve any new AI vendor relationship in the first hundred days outside the existing sanctioned set. The governance framework must be in place before new vendors are added. Pending requests will be deferred or rejected.

I will not deploy any AI workflow in production during the first hundred days. The activity is foundational. The deployment of actual workflows begins in Q1.

I will not announce any of this externally. The work is internal until we have something deployed and operating. Public commitments before deployment create pressure to ship before the work is ready.

How I expect to be held accountable

I propose that the audit committee adopt a standing AI governance agenda item beginning at the next meeting. The item will appear at every meeting through the eighteen-month horizon of this plan. Diana has indicated her willingness to do this. I welcome the visibility because the visibility is what protects the program from drift.

I propose that we set quarterly review points where the progress against the framework, the deployment plan, and the success criteria is examined honestly. Where we are behind, we examine why. Where we are ahead, we examine why as well, because being ahead of a deliberate plan often means we have cut a corner.

I propose that we set one external review point at the twelve-month mark, where an outside advisor with AI and finance expertise reviews the framework, the deployments, and the governance posture. The outside review provides a check on our internal calibration and produces evidence that we will value in Series C diligence.

Closing

Sarah, this memo represents the synthesis of thirty days of conversation and the framework I propose to use as the operating constitution for AI in the Helix finance function. It is conservative in its principles, deliberate in its sequencing, and ambitious in its target outcome. The eighteen-month destination is a finance function that operates with the visibility, the predictability, and the governance maturity that the next investor will expect. The route to that destination is five workflows, deployed in order, governed under a written framework, and owned by named human beings.

I commit to the framework. I commit to the sequence. I commit to the personal accountabilities. I welcome your reactions and your modifications. Whatever we ultimately agree to, I would like it to be in writing, signed by both of us, before deployment begins.

I look forward to your response.

John Campbell



Appendix A · Glossary of Internal Terms

Terms specific to this memo and to the broader AI program that follows. Most are defined in Part 1 of the masterclass; a few are specific to the Helix context.

Architectural framework

The set of principles, structures, and decisions that govern how AI is deployed in the Helix finance function. Drafted in this memo, refined in Part 4 of the masterclass.

Audit trail

An immutable record of every AI workflow invocation, capturing input, model version, output, human review, and timestamps. Required for every workflow under Principle Three.

Cadenza

The AI-powered sales productivity tool currently deployed in the Helix revenue operations function. To be brought onto a corporate master agreement during the foundation period.

Five use cases

The five AI workflows proposed for deployment in this memo. Forecasting Engine, Board Reporting, Pipeline Intelligence, Finance Operations Copilot, Infrastructure and Vendor Intelligence.

Foundation work

The three foundational technical components proposed by Raj Patel and the parallel governance work that must be completed before any workflow can be safely deployed. Scheduled for the quarter preceding workflow deployment.

Governance framework

The written set of policies, processes, and controls that govern AI use across the enterprise. Drafted by the CFO with input from the General Counsel and Head of Security. Subject of Part 4 of the masterclass.

Helix

Helix Cloud Systems, Inc., the company. Used throughout this memo as the working shorthand.

Human-in-the-loop

The architectural pattern in which AI workflows produce output that is reviewed by a named human reviewer before release. Required for all financial outputs under Principle Two.

Operating thesis

The single-sentence articulation of what the finance function must become in order to meet the demands placed on it. Set out in Section 1 of this memo.

Risk tier

A classification of AI use cases by their potential for harm. The classification framework will be specified in the governance framework. Determines the level of review required before deployment.

Shadow AI

AI use across the enterprise that has occurred outside the sanctioned governance framework. Subject to inventory and remediation during the foundation period.

Workflow registry

The central record of every AI workflow deployed in production, including its purpose, owner, governance classification, and operational status. To be built as part of the foundation work.

Appendix B · Assessment

Twenty questions on Part 3. Twelve multiple choice, five short answer, three scenario-based. The answer key with explanations follows.

Part I: Multiple Choice

1. John's operating thesis describes Helix's finance function as primarily facing:

- (a) A revenue growth problem requiring better sales tooling.
- (b) A multi-source signal aggregation problem that the current architecture is not solving.
- (c) A cost discipline problem requiring expense reduction.
- (d) A talent acquisition problem requiring more senior hires.

2. The first of John's seven architectural principles is:

- (a) Human-in-the-loop for all financial outputs.
- (b) Workflows before agents.
- (c) No production AI without an audit trail.
- (d) Every workflow has a named human owner.

3. The readiness principle in Section 4 states that:

- (a) Readiness is the average of all readiness dimensions.
- (b) Readiness is determined by the weakest dimension, not by the average.
- (c) Readiness is determined primarily by data quality.
- (d) Readiness is determined primarily by engineering capacity.

4. The five criteria for evaluating an AI use case candidate, as proposed by John, are:

- (a) Business impact, executive sponsorship, data readiness, governance risk, reversibility.
- (b) Cost, time, scope, quality, risk.
- (c) Strategic fit, financial return, operational feasibility, regulatory clearance, vendor maturity.
- (d) Technical depth, model capability, integration complexity, audit defensibility, user adoption.

5. Which workflow does John propose to deploy first, in Q1 of the next fiscal year?

- (a) The forecasting engine, because it directly addresses the board's predictability concern.
- (b) The pipeline intelligence system, because it has the cleanest data.
- (c) The finance operations copilot, because it builds team capacity and operates on already-clean data.
- (d) The infrastructure and vendor intelligence system, because it is the lowest risk.

6. The forecasting engine is sequenced fourth, in Q4. The principal reason for this placement is:

- (a) Its engineering complexity exceeds the other workflows.
- (b) It depends on the pipeline intelligence workflow as an input and carries the highest failure cost.
- (c) The board has indicated it should be deployed last.
- (d) It requires fine-tuning, which takes longer than the other workflows.

7. The customer health scoring workflow Priya proposed was:

- (a) Adopted as the sixth use case in the program.
- (b) Rejected as out of scope for the finance function.
- (c) Deferred to a second wave because of overlap with the forecasting workflow architecture and capacity constraints.
- (d) Combined directly with the board reporting workflow.

8. John commits to NOT do which of the following in the first hundred days?

- (a) Hire one additional FP&A; analyst.
- (b) Approve any new AI vendor relationship outside the existing sanctioned set.
- (c) Draft the AI governance framework.
- (d) Publish the AI usage policy.

9. The proposed sequencing places foundation work in the quarter before deployments begin because:

- (a) The audit committee has mandated this specific sequence.
- (b) Workflows that depend on the foundation cannot be safely deployed until it exists, and governance must precede deployment under Principle Seven.
- (c) Engineering capacity is unavailable during the deployment quarters.
- (d) Vendor pricing is more favorable for foundation work.

10. John's memo deliberately avoids:

- (a) Naming specific executives by name.
- (b) An executive summary, bullets, or slides.
- (c) References to the board.
- (d) Specific dollar amounts or financial figures.

11. The NDR figure of 118 percent reported by Helix:

- (a) Is correct and matches the calculation method used by the next investor.
- (b) Includes contractual escalators and price increases that, if separated, would produce an expansion-only NDR closer to 113 percent.
- (c) Has already been corrected in the most recent board package.
- (d) Is conservative relative to peer companies of similar size.

12. John proposes to be held accountable through:

- (a) Annual performance reviews only.
- (b) A standing AI governance agenda item at every audit committee meeting, quarterly review points, and a twelve-month external review.
- (c) Monthly briefings to the full board.
- (d) A self-assessment delivered to the CEO at the end of the eighteen-month plan.

Part II: Short Answer

13. In two or three sentences, explain why John structures the memo with the operating thesis in Section 1 rather than placing the use case list at the front. What executive principle does this ordering reflect?

14. John's seventh architectural principle states that governance precedes deployment with no exceptions. In two or three sentences, explain why he insists on no exceptions even for workflows that appear low-risk.

15. The prioritization framework distinguishes between five evaluation criteria and three sequencing thresholds. In two or three sentences, explain the difference between criteria and thresholds and why John treats them separately.

16. John defers the customer health scoring workflow despite the fact that the sixty-percent churn surprise rate is one of the most striking findings from the discovery. In two or three sentences, explain his reasoning and what the deferral signals about how he is making sequencing decisions.

17. In Section 9, John commits explicitly to what he will not do as well as what he will do. In two or three sentences, explain why the negative commitments matter and what they accomplish that the positive commitments alone could not.

Part III: Scenario-Based

18. Scenario: Three weeks after presenting the memo to the CEO, you receive a request from the chief revenue officer to fast-track the pipeline intelligence workflow ahead of the finance operations copilot, on the grounds that the board's predictability concern is most directly tied to pipeline reliability. The CRO offers to fund part of the engineering work from the sales operations budget and proposes a Q1 deployment instead of Q3. In one paragraph of executive prose, describe how you would respond, what you would and would not concede, and what principle from your own memo you would invoke.

19. Scenario: At the audit committee meeting where you present your governance framework, the chair pushes back on Principle Two (human-in-the-loop for all financial outputs) on the grounds that the principle,

as written, may be too restrictive and could prevent automation of high-confidence routine outputs. She asks whether the principle could be modified to allow for exception categories where human review is not required for every individual output. In one paragraph, describe how you would respond, whether you would propose a modification, and what governance protections you would require if the principle were softened.

20. Scenario: Six months into deployment, the finance operations copilot is running in production and the board reporting workflow is in late-stage design. You discover that two of the engineers who built the foundation work have left the company, and the institutional knowledge about how the workflow registry and audit log were architected is now thin. The new engineering hires are competent but have no context. In one paragraph of executive prose, describe what governance and operational steps you would take immediately and what longer-term architectural changes you would propose to reduce the dependency on individual engineering knowledge.

Appendix C · Answer Key with Explanations

Check your answers carefully. The scenario-based questions are deliberately difficult because they ask you to operate as John Campbell would operate, not merely to recall what he wrote.

Multiple Choice Answers

Question 1: (b)

John's operating thesis frames the finance function as facing a multi-source signal aggregation problem. The forecast misses, churn surprises, and infrastructure waste are all surfaces of the same underlying condition: fragments of signal generated by competent people in different functions, never aggregated into a coherent view. See Section 1.

Question 2: (b)

The first principle is workflows before agents. The principle establishes the conservative default architecture for the eighteen-month horizon, with autonomous agents permitted only with specific justification approved by the CFO and the audit committee chair. See Section 5.

Question 3: (b)

The readiness principle states that the weakest dimension dominates. A company with strong data and weak governance is not ready for serious AI deployment in finance, because the governance gap will produce the failure that ends the program. See Section 4.

Question 4: (a)

Business impact, executive sponsorship, data readiness, governance risk, and reversibility are the five criteria. Each is a question asked of every candidate workflow. A workflow that fails on any one is either not yet ready or not a good candidate at this stage. See Section 6.

Question 5: (c)

The finance operations copilot is sequenced first. It produces immediate visible value to the finance team itself, gives David and Elena first experience operating under governance, operates on clean NetSuite data, and deploys first into the function that owns the governance framework. See Section 8.

Question 6: (b)

The forecasting engine is sequenced fourth because it depends on the pipeline intelligence workflow as an input, requires the team to have experience operating multiple workflows under governance, and carries the highest failure cost. Deploying the highest-stakes workflow after three lower-stakes workflows is the right risk posture. See Section 8.

Question 7: (c)

The customer health scoring workflow was deferred to a second wave because of architectural overlap with the forecasting workflow, heavier data preparation requirements, and customer success function bandwidth constraints. The capability will be partially captured as a side effect of the forecasting workflow. See Section 7.

Question 8: (b)

John commits explicitly NOT to approve any new AI vendor relationship outside the existing sanctioned set during the first hundred days. The governance framework must be in place before new vendors are added. See Section 9.

Question 9: (b)

Foundation work precedes deployment because workflows that depend on the foundation cannot be safely deployed until it exists, and because Principle Seven establishes that governance precedes deployment uniformly. See Sections 5 and 8.

Question 10: (b)

John deliberately avoids an executive summary, bullets, or slides. He has written the memo at the length and in the form Sarah requested, with prose that reveals the reasoning rather than summarizing the conclusions. See the memo header and prologue.

Question 11: (b)

Reported NDR of 118 percent includes contractual escalators and price increases that, if separated from true expansion bookings, produce an expansion-only NDR closer to 113 percent. The correction must be made before the next investor diligence. See Section 3.

Question 12: (b)

John proposes a standing audit committee agenda item, quarterly review points, and a twelve-month external review. The triple structure protects the program from drift, internal complacency, and miscalibration against external standards. See Section 9.

Short Answer Explanations

13. Thesis first, use cases later

John structures the memo with the operating thesis in Section 1 because the use cases must be derived from the thesis rather than the thesis being derived from the use cases. If the memo leads with the five use cases, the reader is being asked to accept the conclusions before seeing the reasoning. By leading with the thesis, John forces the reader to engage with the diagnosis before encountering the prescription. This ordering reflects the executive principle that the conclusion of an argument should be the last thing the reader encounters in a serious memo, not the first thing, because serious memos are about how the writer arrived at the conclusion rather than merely what the conclusion is.

14. No exceptions to governance precedence

John insists on no exceptions because exceptions undermine the framework over time. Once one workflow has been deployed without governance precedence on the grounds that it appeared low-risk, the next workflow argues for the same exception, and the framework becomes a set of guidelines rather than a discipline. The integrity of the governance framework comes from its uniform application, not from its specific rules. The cost of slowing some low-risk deployments by a few weeks is far smaller than the cost of erosion in the framework's authority over the program's eighteen-month horizon.

15. Criteria versus thresholds

The five criteria are tests of whether a workflow is a good candidate at all. A workflow that fails on any one of the five is either not yet ready for deployment or is the wrong workflow for the program at this stage. The three thresholds are tests of when, among the qualifying workflows, a particular workflow should be deployed. Criteria filter the candidate pool. Thresholds sequence the qualifying candidates. John treats them separately because they answer different questions. Conflating them would make every deployment decision a single judgment call rather than a structured two-step process, which would reduce the discipline and the defensibility of the sequencing.

16. Deferring the customer health workflow

John defers the customer health workflow despite the striking churn surprise rate because deferral reflects honest sequencing rather than alphabetical sequencing. The architecture overlaps with the forecasting workflow, the customer success function does not have the bandwidth to lead the design, and the capability will be partially captured as a side effect of the forecasting workflow. The deferral signals that John is making sequencing decisions based on architectural coherence and operational capacity rather than on the relative urgency of individual pain points. It also signals that he is willing to accept some operational pain in the near term in order to build the right architecture rather than the most immediate response.

17. Why the negative commitments matter

The negative commitments establish what John will refuse to do, which is harder and more telling than what he will do. A list of positive commitments is what most executives produce in a memo of this kind, and the list is often inflated because it costs nothing to add another item. The negative commitments cost something: they impose discipline on the writer's own future choices. By committing not to approve new vendors, not to deploy workflows, and not to make external announcements during the first hundred days, John is preempting the temptations that will arise during that period and binding his own future self to the discipline of the framework. The negative commitments make the framework durable in a way the positive commitments alone cannot.

Scenario Discussions

18. The CRO's fast-track request

The right response is to decline the fast-track while acknowledging the legitimate concern that underlies it. The CRO is correct that pipeline reliability is closely linked to the board's predictability concern. The CRO is also asking for a sequencing change that violates two principles of the framework. First, the foundation work for the pipeline intelligence workflow has not yet been completed; the workflow depends on the MCP server in front of the warehouse and on the workflow registry, neither of which will be operational until the end of the foundation quarter. Second, the team has no experience operating an AI workflow under governance, and deploying first into a function outside finance, on the strength of a budget contribution from sales operations, sets the precedent that governance can be traded for budget. The principle to invoke is Principle Seven: governance precedes deployment, with no exceptions. The constructive response is to engage the CRO on what we can do now: tighten the rep stage discipline that Lisa Mahoney can implement immediately, surface the stage-five slip rate data over the eight-quarter window, and begin parallel design work on the pipeline intelligence workflow so that it is ready for Q3 deployment as planned. The CRO's budget contribution can be welcomed as an investment in the foundation work that will benefit his function as much as anyone else's.

19. The audit committee chair's push on human-in-the-loop

The right response is to engage seriously with the concern rather than to defend the principle as written. The chair is making a substantive point. A principle that requires human review of every individual output, regardless of confidence level or risk category, may produce more friction than control. The constructive modification would be to specify that human review is required at the workflow level rather than at the individual output level, with two protections. First, every workflow must define its own human review pattern as part of its design, including which outputs require individual review and which can be batched or sampled, and the design must be approved before deployment. Second, no workflow can move from individual review to batched or sampled review without a documented performance history demonstrating that the output quality is consistent enough to justify the change. The principle becomes: human-in-the-loop architecture for every workflow, with the specific review cadence determined by the workflow's risk tier and demonstrated reliability. This preserves the discipline while permitting the operational maturity that the chair is rightly insisting on as a path forward. The modification should be documented in the framework with the chair's endorsement and revisited annually.

20. The departure of foundational engineers

The immediate steps are operational. First, conduct a knowledge audit of the workflow registry and audit log within two weeks, with the new engineering hires and any remaining team members documenting every architectural decision and every operational pattern that has been institutional knowledge until now. The audit should produce a written reference document that any future engineer can use to operate and modify the foundation systems. Second, suspend any deployment work that depends on the foundation systems until the audit is complete and the new engineers have demonstrated they can operate the systems independently. Third, brief the audit committee chair on the situation; the chair should not learn of the dependency risk after a failure rather than before. The longer-term architectural changes are more important. The dependency on individual engineering knowledge is itself a control weakness, and the framework should require that every component of the foundation be documented sufficiently for any qualified engineer to operate it. This is a standing requirement, not a one-time response to the current situation. The deeper lesson is that foundational components in a finance AI program must be treated with the same documentation discipline that the financial reporting systems themselves require. The fragility we have just discovered should be the last time it is permitted.



End of Part 3

The Memo

John has now produced the memo. The framework is on paper. The five use cases are named, justified, and sequenced. The principles are declared. The personal commitments are recorded. The work of the next eighteen months has its constitution.

In Part 4, the masterclass turns to the governance framework itself. John has committed to drafting it within the first thirty days following acceptance of the memo. Part 4 is the full treatment of what that framework must contain: the three pillars of AI governance, the human-in-the-loop architecture, the audit trail requirements, the risk classification system, the vendor risk framework, the methodology for embedding governance into workflow design, the regulatory and external context, and John's governance charter as the written artifact that becomes the operating constitution.

Part 4 is the longest of the foundational parts and the most rigorous. It is also the part that, once internalized, will transform how you read the use case parts that follow. Every use case in Parts 5 through 9 will reference the governance methodology Part 4 establishes. Every architectural decision will be evaluated against the principles Part 4 codifies. The masterclass turns now from listening and synthesis to the construction of the framework that will govern everything that follows.

