Highspring | vaco BY HIGHSPRING

AGILITY AT WORK.

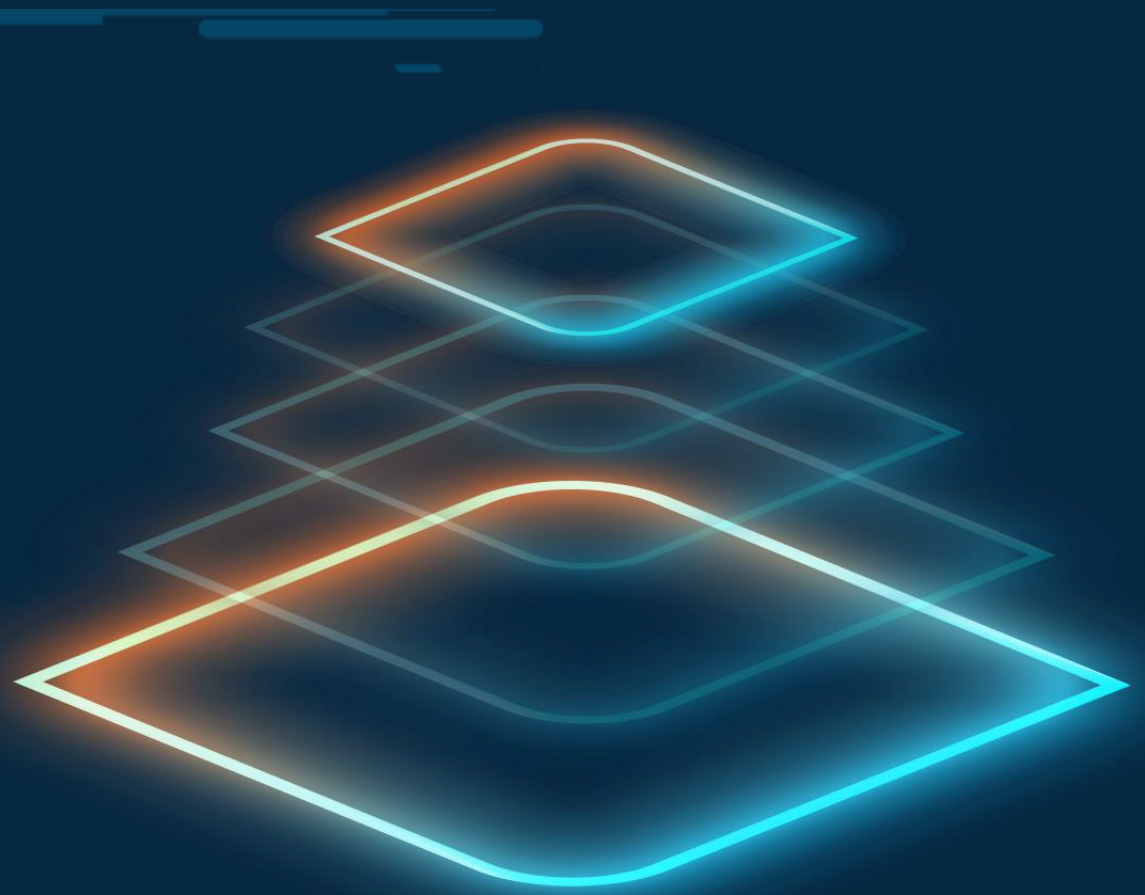

CPE LUNCH & LEARN

AI in Financial Reporting

Practical Governance & Controls

August 19, 2026

Today’s Agenda

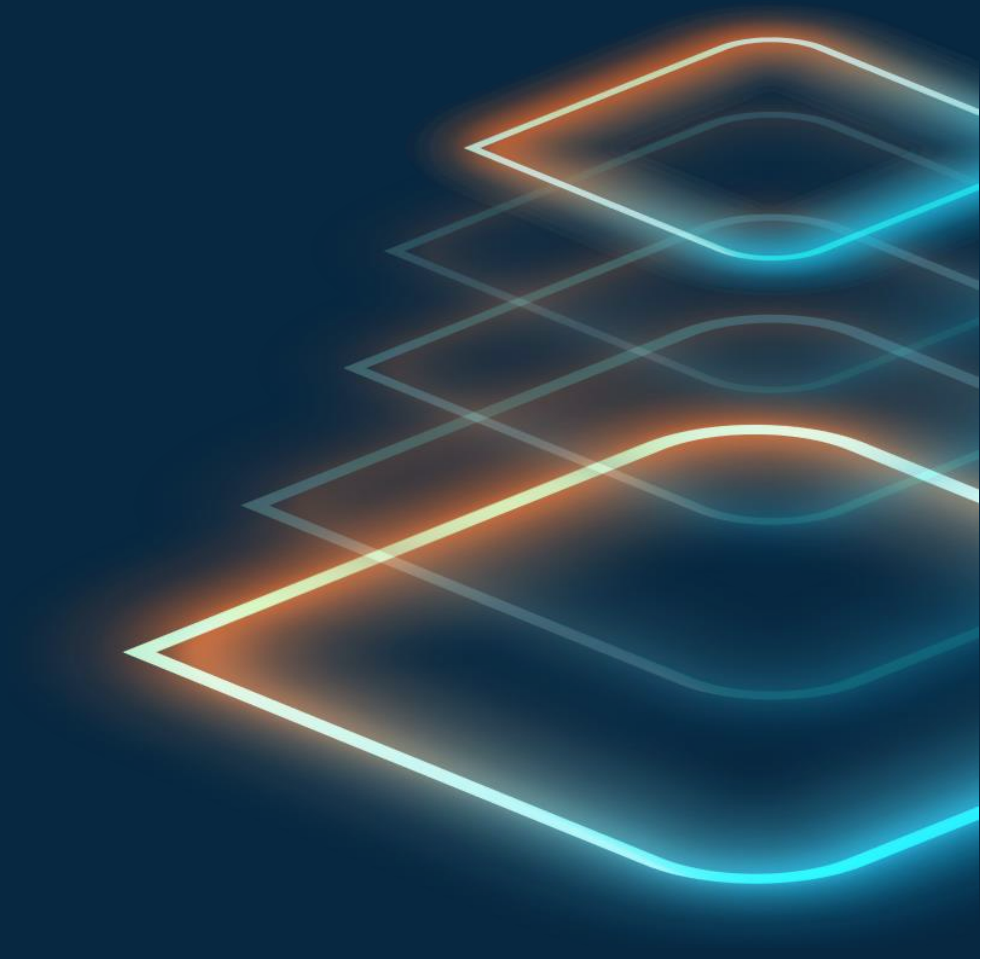

12:15 – 12:20 PM	**Welcome & why this matters to finance**
12:20 – 12:28 PM	**Part 1 — AI readiness** Strategy, data integrity, and reporting foundations
12:28 – 12:37 PM	**Part 2 — AI governance** ISO/IEC 42001, oversight, and policy development
12:37 – 12:41 PM	**Part 2 (continued)** Standing up the AI governance board
12:41 – 12:55 PM	**Part 3 — AI in the reporting process** Scoping AI into ICFR and SOX, risk assessment, and controls
12:55 – 1:05 PM	**Part 4 — Financial reporting use cases** Close, reconciliations, IPE, and disclosure narratives Controls and evidence that keep reported numbers audit-ready
1:05 – 1:08 PM	**Closing — the 90-day roadmap**
1:08 – 1:10 PM	**Q&A**

Your Speaker Today

With deep expertise across SOX, financial reporting controls, and cybersecurity, Christian brings a practitioner's perspective grounded in real-world implementation across highly regulated industries.

CHRISTIAN HEFFRON

Highspring Consulting Partner, Risk & Regulatory Leader

Christian has extensive experience as a Chief Audit Executive leading outsourced and co-sourced internal audit functions. He also leads numerous SOX implementations as well as SOX compliance services. Prior to joining Highspring, Christian served as the risk advisory partner for a mid-sized CPA firm, leading the firm's internal audit and risk management practice. Christian also led internal audit functions at Freddie Mac, TimeWarner and Altria. Christian started his career with PwC and is a Chartered Accountant (CA) and Certified Information Systems Auditor (CISA).

Part 1

AI READINESS

Why This Matters to Finance

AI is no longer a future consideration for finance, it's already here. The question is not whether to engage, but how to govern and control it responsibly.

FINANCIAL CLOSE
AI-assisted close processes and period-end procedures

FINANCIAL REPORTING
Automated narrative generation and disclosure drafting

RECONCILIATIONS
Intelligent matching and exception identification

JOURNAL ENTRY REVIEW
Anomaly detection and risk-based sampling

CONTROLS AND EVIDENCE
Automated evidence collection supporting key reporting controls

IPE VALIDATION
Completeness and accuracy checks on information produced by the entity

Finance leaders need confidence that AI outputs touching the financial statements are governed, controlled, and audit-ready.

Foundational data maturity is essential

95% of organizations aren’t getting a return on their AI investments

Source: State of AI in Business 2025 Report

Finance leaders face a talent and AI gap

42% of finance leaders are confident in hiring digital finance talent

Source: Gartner

Tech leaders must adapt quickly

By 2030, 75% of IT work will be humans augmented by AI—and 25% done by AI alone

Source: Gartner

Setting the Stage

Four key drivers of organizational AI readiness guide the design and implementation of an effective AI program.

BUSINESS STRATEGY

How AI supports business goals and value creation

TECHNOLOGY & DATA STRATEGY

Data and infrastructure to fuel and run AI at scale

AI GOVERNANCE

Processes, controls, and accountability for responsible AI use

ORGANIZATION & CULTURE

Leadership vision, operating model, skills, and culture

Adapted from Highspring's AI Readiness Framework

AI readiness is a program, not a tool

Successful organizations approach AI readiness holistically, addressing strategy, data, technology, risk, talent, and change together. When these elements are coordinated, AI initiatives scale with confidence and deliver sustained value.

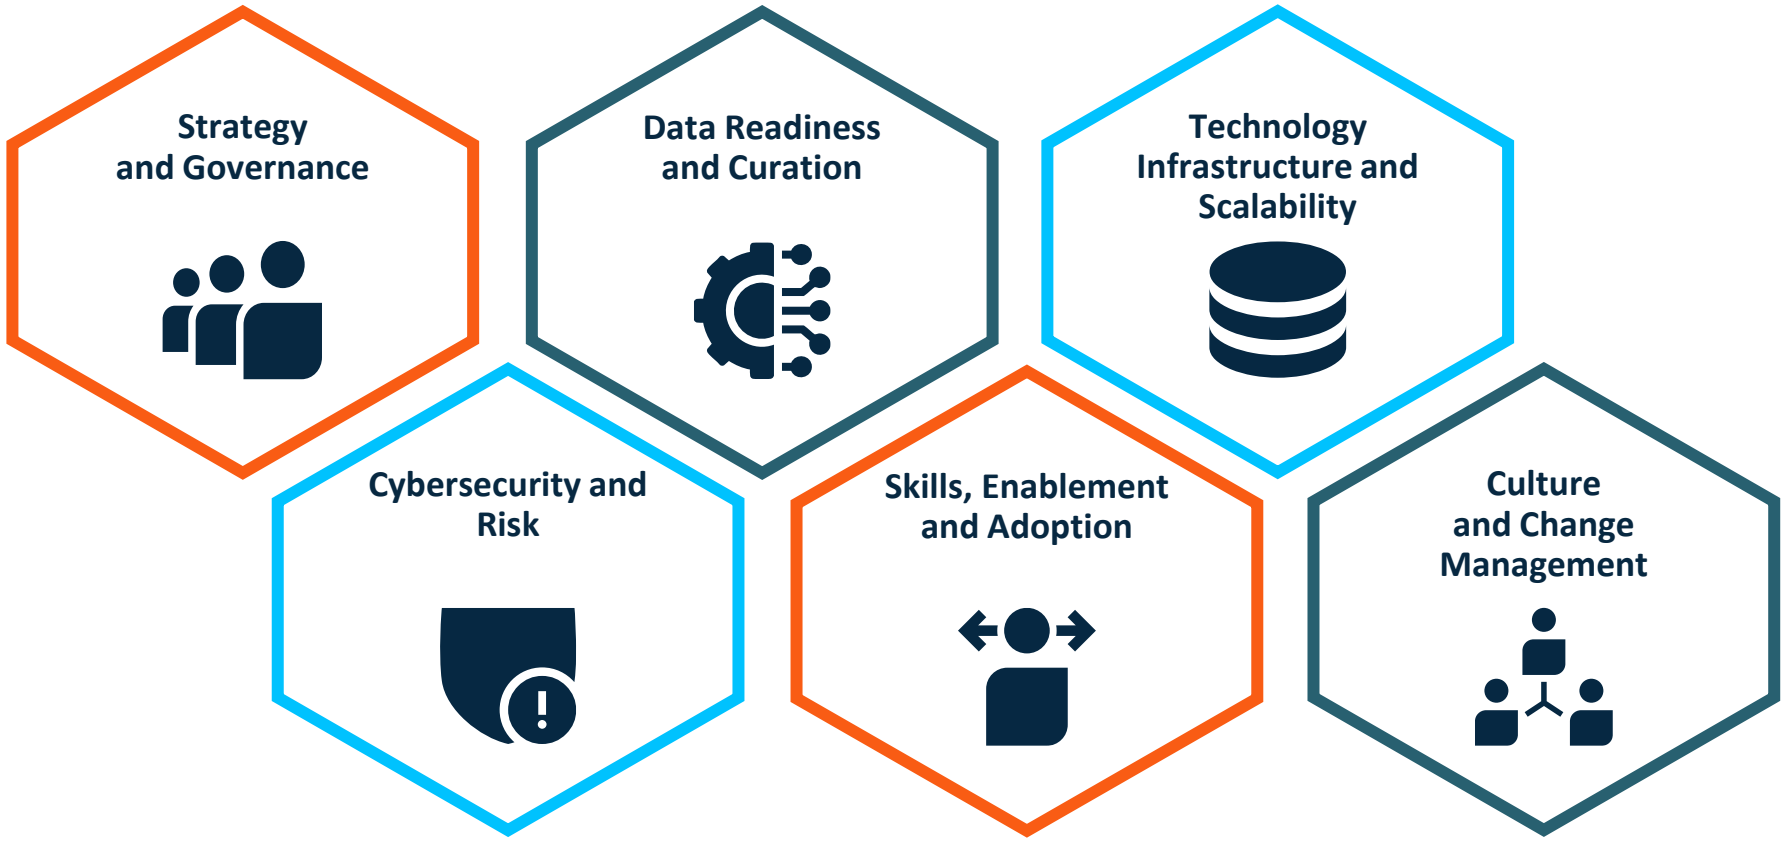

The roadmap to AI maturity

Build a strong data foundation before scaling AI or business transformation.

These six keys turn data into actionable insights.

Data & AI readiness maturity rubric

Below is the rubric based on the Strategic Data Management & Analytics (SDM&A) Framework used to evaluate a company's data maturity in the defined data-related categories.

Maturity levels	Data governance	Reporting and analytics	Data delivery	Data architecture	Technology and operations	Metadata	Data quality and metrics	Data science
Strategic (Level 5)	Enterprise data governance operational	Data marketplace operational to perform advanced reporting	Fully supported and operational with real time and batch	LOBs adhere to and Endorse enterprise DA and models	Highly functional production BI releases on a frequent basis	Data lineage tools and procedures are followed	DQ program has been established	Enterprise model governance and reuse
Optimized (Level 4)	Data governance in place – not all LOBs	Enterprise reporting tool allows self-serve and operational reporting	Emerging technologies being introduced	Enterprise data models drive decisions	Change management process which BI solutions adhere to	Metadata repositories are maintained	LOBs working on DQ improvements at source	Standardized templates and validation
Defined (Level 3)	Governance and stewardship recognized	Key reports delivered in batch in dashboards	Ingestion and ETL standards and framework	Data warehouses and data lakes in place	Shared resources exists to provide support, but not fully utilized	Lineage is tracked manually via SQL	DQ being handled at ETL / ingestion layer	Active modeling, reuse awareness
Managed (Level 2)	Directors and management advocating	Key reports delivered in batch in excel	A common ETL platform is used (not by all)	Applications drive data arch. and models	Central IT and shared services providing minimal support	ETL source to target mappings exist	DQ handled within excel and reporting layer	Isolated models, limited standards
Ad Hoc (Level 1)	SMEs act as "stewards"	Reporting handled by individuals in multiple tools	Combination of ingestion techniques (tools, SQL)	No enterprise data models or standards	Each LOB BI developers providing support	Data exists in multiple silos (Excel, Access)	Improving DQ not recognized	Siloed, ungoverned modeling

Foundational capabilities for AI readiness

Common problems	Key capabilities	Business outcomes
"I don't have an accurate, 360-view of the key data needed, nor do I know which definitions to use to make decisions."	**Governance and master data management (MDM)**	**Empower teams to "speak the same language"** • Actively manage data lineage and critical data elements • Establish a scalable data environment and observability
"We have so many different systems, but data across them all is often inconsistent, missing, or incorrect."	**Enriched data quality**	**Accurate, conforming, and timely data—no rework** • Build and continuously monitor data quality dashboards • Data quality issue resolution framework
"Our data is saved in various Excel files and operational systems, so each team relies on their own different 'source of truth' for reporting purposes."	**Data centralization (data warehouse / lake / lakehouse)**	**Unify data – One version of truth** • Connect disparate source systems to target system • Drive reporting off data warehouse, not operational systems
"My team spends weeks manually wrangling data in Excel. It takes us so long to pull together data for our reports that we don't even have time to analyze and act on our findings."	**Process optimization, automation and AI**	**Less time wrangling, and more time analyzing** • Streamline and optimize highly manual processes for automation • Develop reusable assets
"We have 100+ reports, but don't even know which metrics are correct to use for decision making."	**Proactive analytics and insights (metrics that matter)**	**Drive your business with data** • Streamline reporting and real-time BI / visualization capabilities • Develop predictive analytics and insights
"I'm sure there are ways to help me get out of Excel and better use the data to work more efficiently, but what does that look like?"	**Data literacy and culture**	**Data democratization** • Establish clear policy and procedures and roles and responsibilities • Self-service analytics, and citizen development

What does success look like?

Data strategy and governance should be viewed as a continuous operating model rather than a one-time project or program. Improvements come from enhanced operational efficiency, lower risk and compliance expenses, and greater confidence in data for informed decision-making. **Enterprise data programs have the potential to increase ROI by as much as 300% within three years, with a positive return on investment often realized in the first year.**

Key learnings from prior enterprise data engagements

Long Term Sustainability vs. Short Term Considerations	**Investing in scalable and adaptable solutions,** and avoiding "patches and manual workarounds" is paramount to **sustainable growth, specifically as an organization begins to implement AI and ML**
Enterprise Tool Standardization vs. Function Specific Tools	Utilizing **tools that can be flexed across the enterprise, including AI enabled tools,** instead of niche tools for each team, creates standardization, improved data quality and insights, leading to **better business outcomes**
Frictionless Data Governance vs. Team Specific Data Practices	**Creating buy-in and** ensuring that governance protocols meet use case requirements, allows for **an agile and robust governance model with the benefits of reliable data, without stifling the business** with rigid standards and heavy admin burdens

Typical engagement outcomes

50%+	**Efficiency Gains** from Manual Governance Process Elimination
15%+	**Technology Cost Savings** from Data & System Rationalization Initiatives
75%+	**Reporting Speed Improvement** from Data Governance & Reporting Initiatives
19x	**Data-driven companies** are upwards of 19 times more likely to be profitable

Part 2

AI GOVERNANCE

PART 2: AI GOVERNANCE

Why do we need an AI governance framework?

Artificial intelligence is increasingly embedded in critical business processes, decision-making systems, and customer-facing services. While AI offers significant opportunities for innovation, efficiency, and competitive advantage, it also introduces **new categories of risk** related to ethics, safety, transparency, security, regulatory compliance, and societal impact. Without a structured approach to governing AI, organizations face heightened exposure to operational failures, legal liabilities, reputational damage, and loss of stakeholder trust.

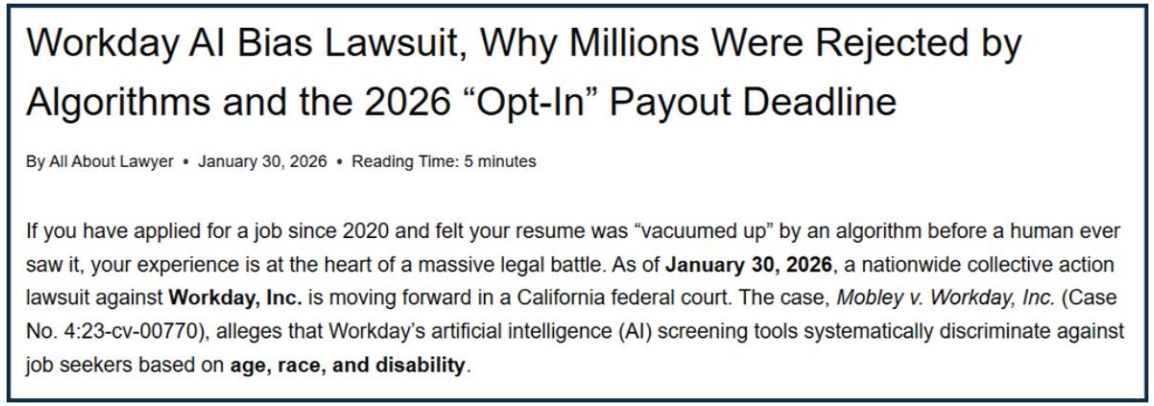

Workday AI Bias Lawsuit, Why Millions Were Rejected by Algorithms and the 2026 “Opt-In” Payout Deadline

By All About Lawyer • January 30, 2026 • Reading Time: 5 minutes

If you have applied for a job since 2020 and felt your resume was “vacuumed up” by an algorithm before a human ever saw it, your experience is at the heart of a massive legal battle. As of **January 30, 2026**, a nationwide collective action lawsuit against **Workday, Inc.** is moving forward in a California federal court. The case, *Mobley v. Workday, Inc.* (Case No. 4:23-cv-00770), alleges that Workday’s artificial intelligence (AI) screening tools systematically discriminate against job seekers based on **age, race, and disability**.

Home / Technology / Tech News / Deloitte's AI fiasco: Why chatbots hallucinate and who else got caught

Deloitte's AI fiasco: Why chatbots hallucinate and who else got caught

As Deloitte returns part of its fee for a flawed AI-assisted government report in Australia, the case spotlights how generative AI fabricates facts and how it's not the first time this has happened

The Washington Post
Democracy Dies in Darkness

Resubscribe joseantonio.ortega

Business Economy Economic Policy Personal Finance Work Technology Business of Climate

EXCLUSIVE

Arrested by AI: Police ignore standards after facial recognition matches

Confident in unproven facial recognition technology, sometimes investigators skip steps; at least eight Americans have been wrongfully arrested.

NBC NEWS OLYMPICS POLITICS U.S. NEWS WORLD HEALTH SPORTS LOCAL SHOPPING WATCH SUBSCRIBE

ARTIFICIAL INTELLIGENCE

Leading AI companies' safety practices are falling short, new report says

Independent experts said in the report that powerful AI systems face critical gaps in safety protocols even though capabilities continue to increase.

ISO/IEC 42001 — A Starting Point for AI Governance

FOUR PILLARS

AI MANAGEMENT SYSTEMS

Structured framework to oversee and manage AI implementation, operation, and risks

AI RISK ASSESSMENT

Systematic approach to identifying and mitigating risks throughout the AI lifecycle

AI IMPACT ASSESSMENT

Evaluation of the impact of AI on individuals and society

DATA PROTECTION AND AI SECURITY

Complying with privacy laws and securing AI systems against threats

AI LIFECYCLE (ISO 22989)

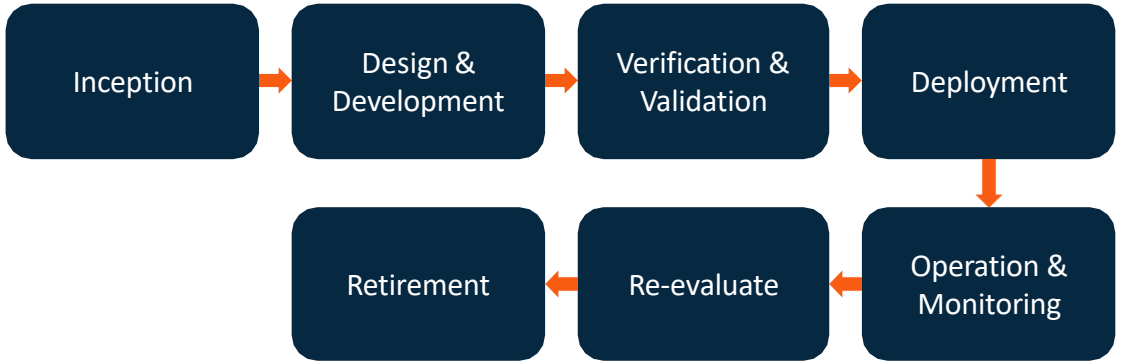

ISO/IEC 42001 provides a management system framework covering the full AI lifecycle, addressing ethics, accountability, transparency, and data privacy.

Governance and security | How they integrate

AI Governance encompasses a collaborative framework where multiple components work in harmony to ensure the development and deployment of artificial intelligence occur within safe, ethical, and accountable boundaries. By integrating policies, technical safeguards, and continuous monitoring, AI Governance creates a robust environment that promotes responsible innovation while mitigating risks associated with AI technologies.

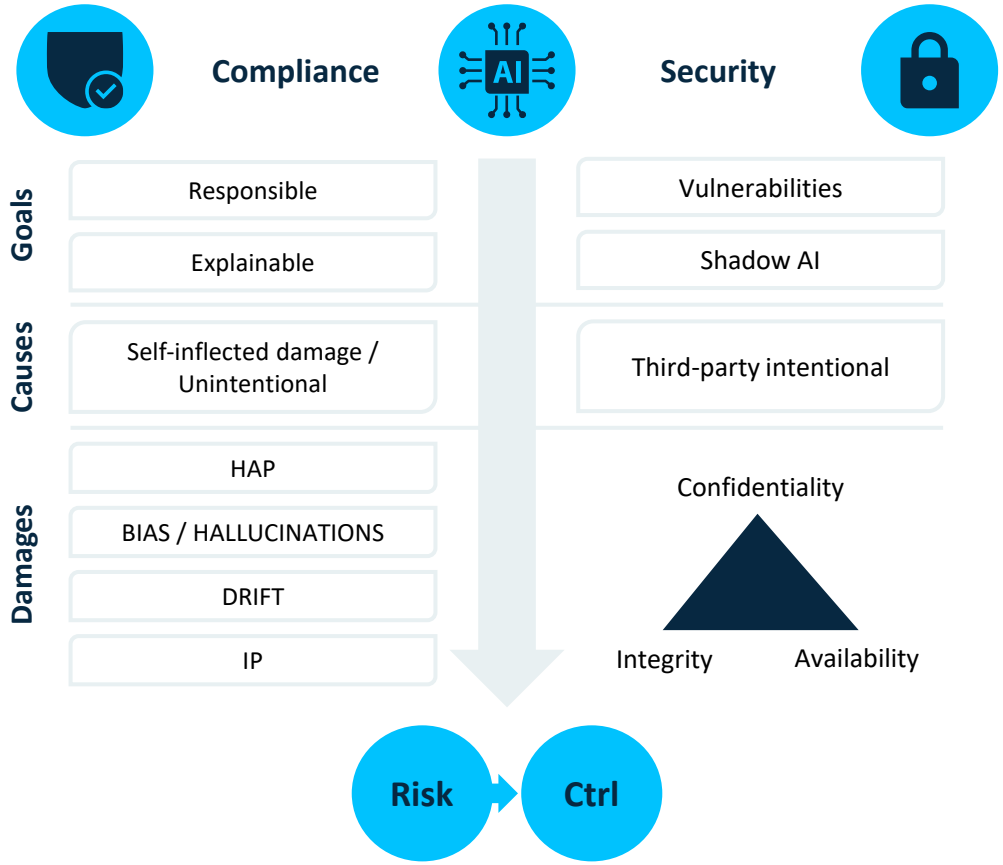

Layers of protection

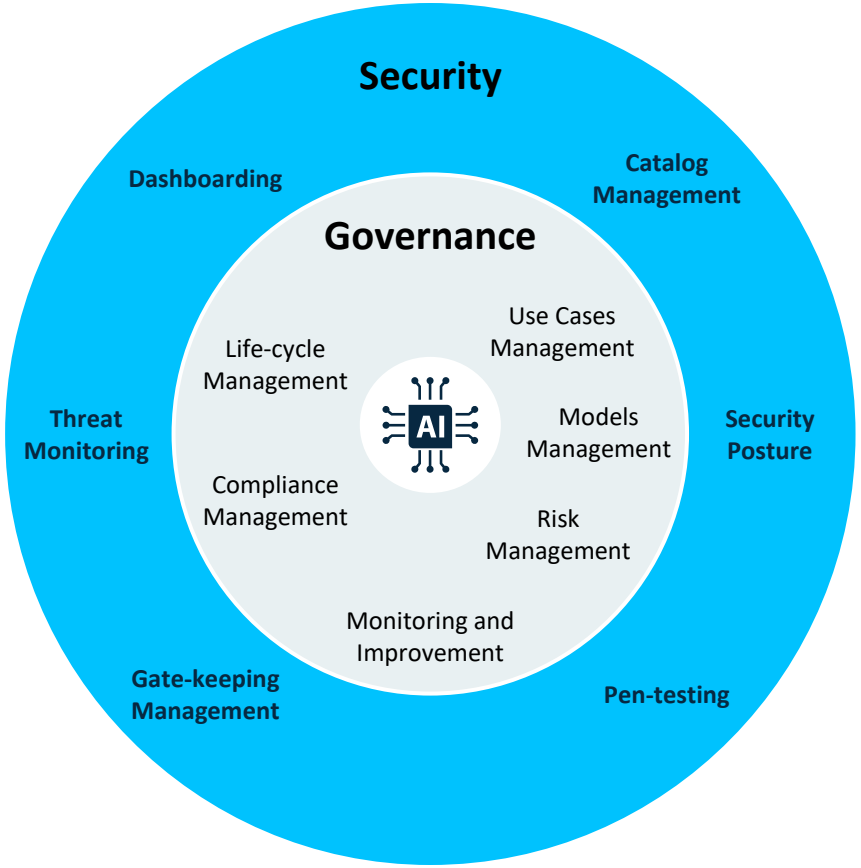

AI Governance Enablement In-practice

EXECUTIVES

Need to Know:

- Enterprise AI risk tiers and approval thresholds
- Decision & escalation triggers for high-impact AI use cases
- Where AI use intersects with ICFR / regulatory exposure
- Oversight dashboards & exception reporting signals

MODEL OWNERS

Need to Know:

- Intake, impact screening, and re-approval triggers
- Fit-for-purpose attestation responsibilities
- Monitoring & human-in-loop requirements by risk tier
- Documentation and evidence expectations

BUILDERS (IT / Data / Engineers)

Need to Know:

- Approved model patterns and technical guardrails
- Data usage, privacy, and access control requirements
- Validation & testing depth by risk tier
- Change management and version traceability standards

BUSINESS USERS

Need to Know:

- Acceptable use and restricted AI use cases
- When human review is required before action
- Prompting & output-reliance do’s and don’ts
- How to report anomalies or unexpected outputs

KEY PRINCIPLE: Most AI risk comes from accidental misuse, not malice

Common AI Governance Operating Models

Option 1: Centralized (The "Fortress" Model)

Core Structure
- Central AI / Risk office owns approvals
- One intake and review for all use cases
- Uniform controls and standards

Strengths
- Maximum control and consistency
- Strong regulatory defensibility
- Clear audit trail

Tradeoffs
- Slower decision velocity
- Can become a bottleneck
- Adoption may bypass the process

Best Fit When
- High regulatory sensitivity
- AI touches financial reporting
- Low risk tolerance

Option 2: Federated (The "Hub & Spoke Model")

Core Structure
- Central standards and guardrails
- Business units own use cases
- Council reviews high-risk only

Strengths
- Balances control with speed
- Scales across functions
- Preserves oversight

Tradeoffs
- Central policy and risk tiers
- Business owns execution
- Council approves high impact

Best Fit When
- Multiple AI initiatives across functions
- Mix of low- and higher-risk uses

Option 3: Decentralized (The "Agile" Model)

Core Structure
- Business teams govern locally
- Minimal central approval
- Lightweight central guidance

Strengths
- Fast experimentation
- Low process friction
- Strong local ownership

Tradeoffs
- Higher shadow AI risk
- Inconsistent controls
- Limited risk visibility

Best Fit When
- Low-risk productivity use cases
- Mature risk culture in teams
- Strong monitoring already exists

KEY PRINCIPLE: Clarity of decision rights matters more than org chart.

AI operating model

Effective AI governance sets the parameters for AI use across the enterprise, resolves issues and escalations, and enables the organization to deploy AI with confidence in its data and controls.

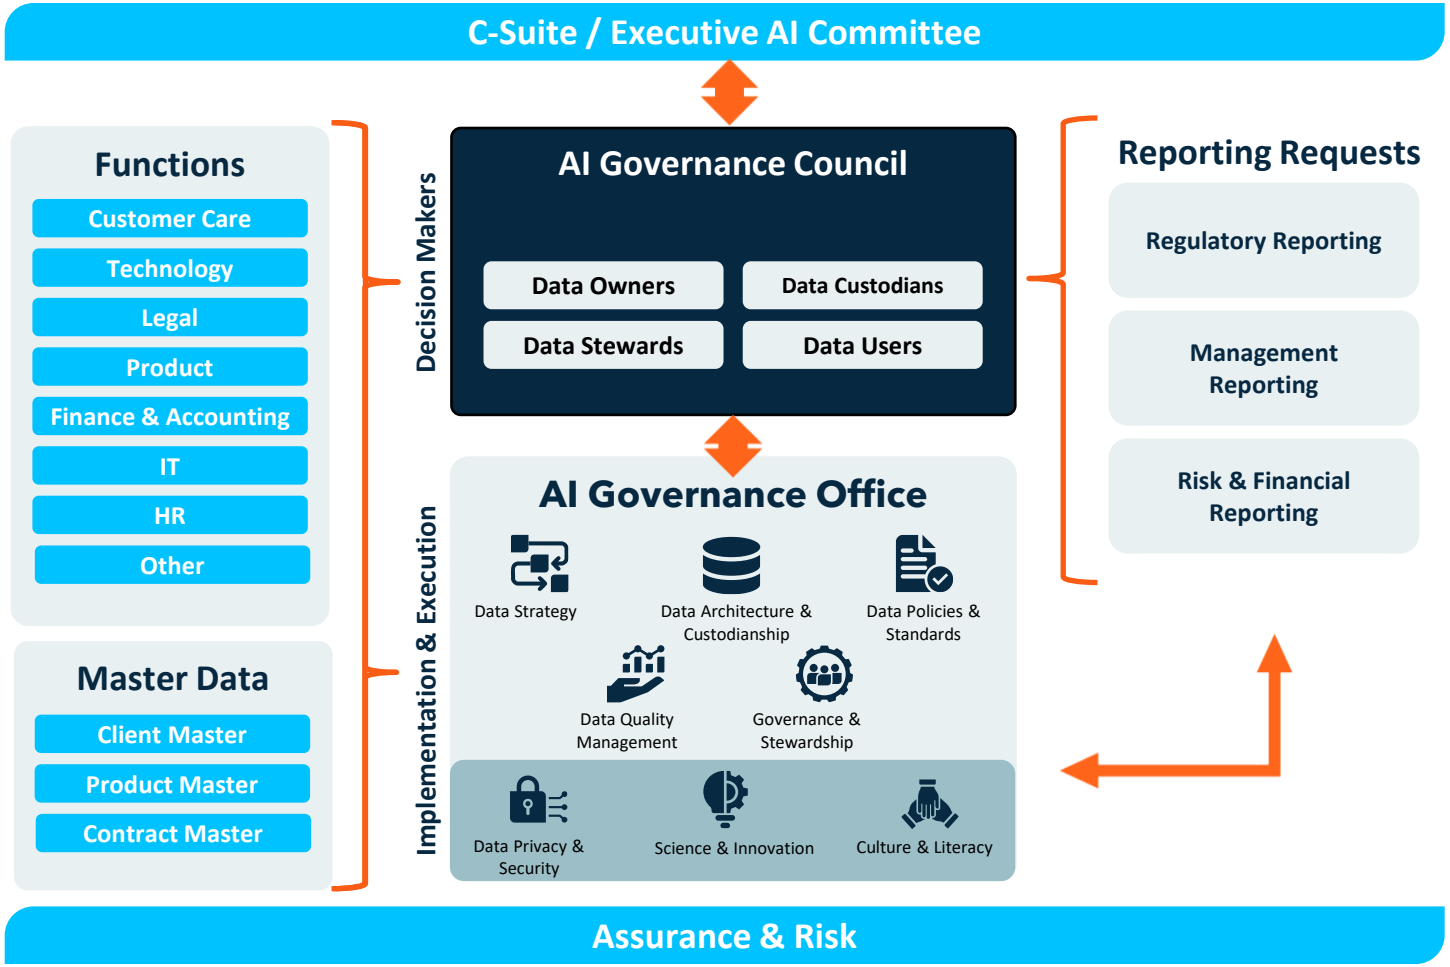

Executive AI Committee
- Accountable for the data strategy
- Represented by the Board and AI Governance Council leadership

AI Governance Council
- Responsible for defining the data strategy, focused on data customer use-case requirements
- Accountable for executing the data strategy

AI Governance Office
- Responsible for executing the data strategy
- AI Governance Council Chair

Functions
- Customers for operational data and creators of reporting inputs, drive governance requirements
- Global and regional first-line and second-line of defense functions with data requirements

Reporting Requirements
- Change the Engine and Run the Engine initiatives feeding in data requirements

Assurance & Risk
- Provides independent assurance over data to senior management and the board

AI Governance Roles – Accountability and Segregation of Duties

ROLE	PRIMARY DUTIES	KEY RESPONSIBILITY
GOVERNANCE COUNCIL	• Sets enterprise AI governance standards & risk thresholds • Reviews and approves high-impact / high-risk AI use cases • Establishes escalation and exception decision paths	• Confirms risk tiering and required control level • Arbitrates cross-functional ownership disputes • Triggers re-review when material change events occur • Provides executive oversight reporting to leadership
BUSINESS OWNER (The Customer)	• Defines business purpose and decision scope of AI use case • Confirms acceptable use boundaries and limitations • Ensures outputs are reviewed before critical decisions	• Attests model output is fit-for-purpose • Owns downstream decision impact • Ensures human-in-the-loop where required by risk tier • Initiates reassessment when use case scope changes
IT / DATA / ENGINEERING	• Designs and configures model architecture and pipelines • Implements technical guardrails and data controls • Maintains versioning and change traceability	• Enforces data source controls and access restrictions • Implements monitoring for drift and anomalies • Executes model updates through controlled change process • Maintains technical documentation and lineage records
RISK / COMPLIANCE	• Performs independent risk and impact assessments • Defines validation and testing requirements by risk tier • Maintains AI control framework and validation criteria	• Conducts independent challenge and model review • Validates controls before high-risk deployment • Reviews monitoring results and exception trends • Recommends approval, restriction, or remediation actions

KEY PRINCIPLE: Builders design — Risk validates — Business owns outcome.

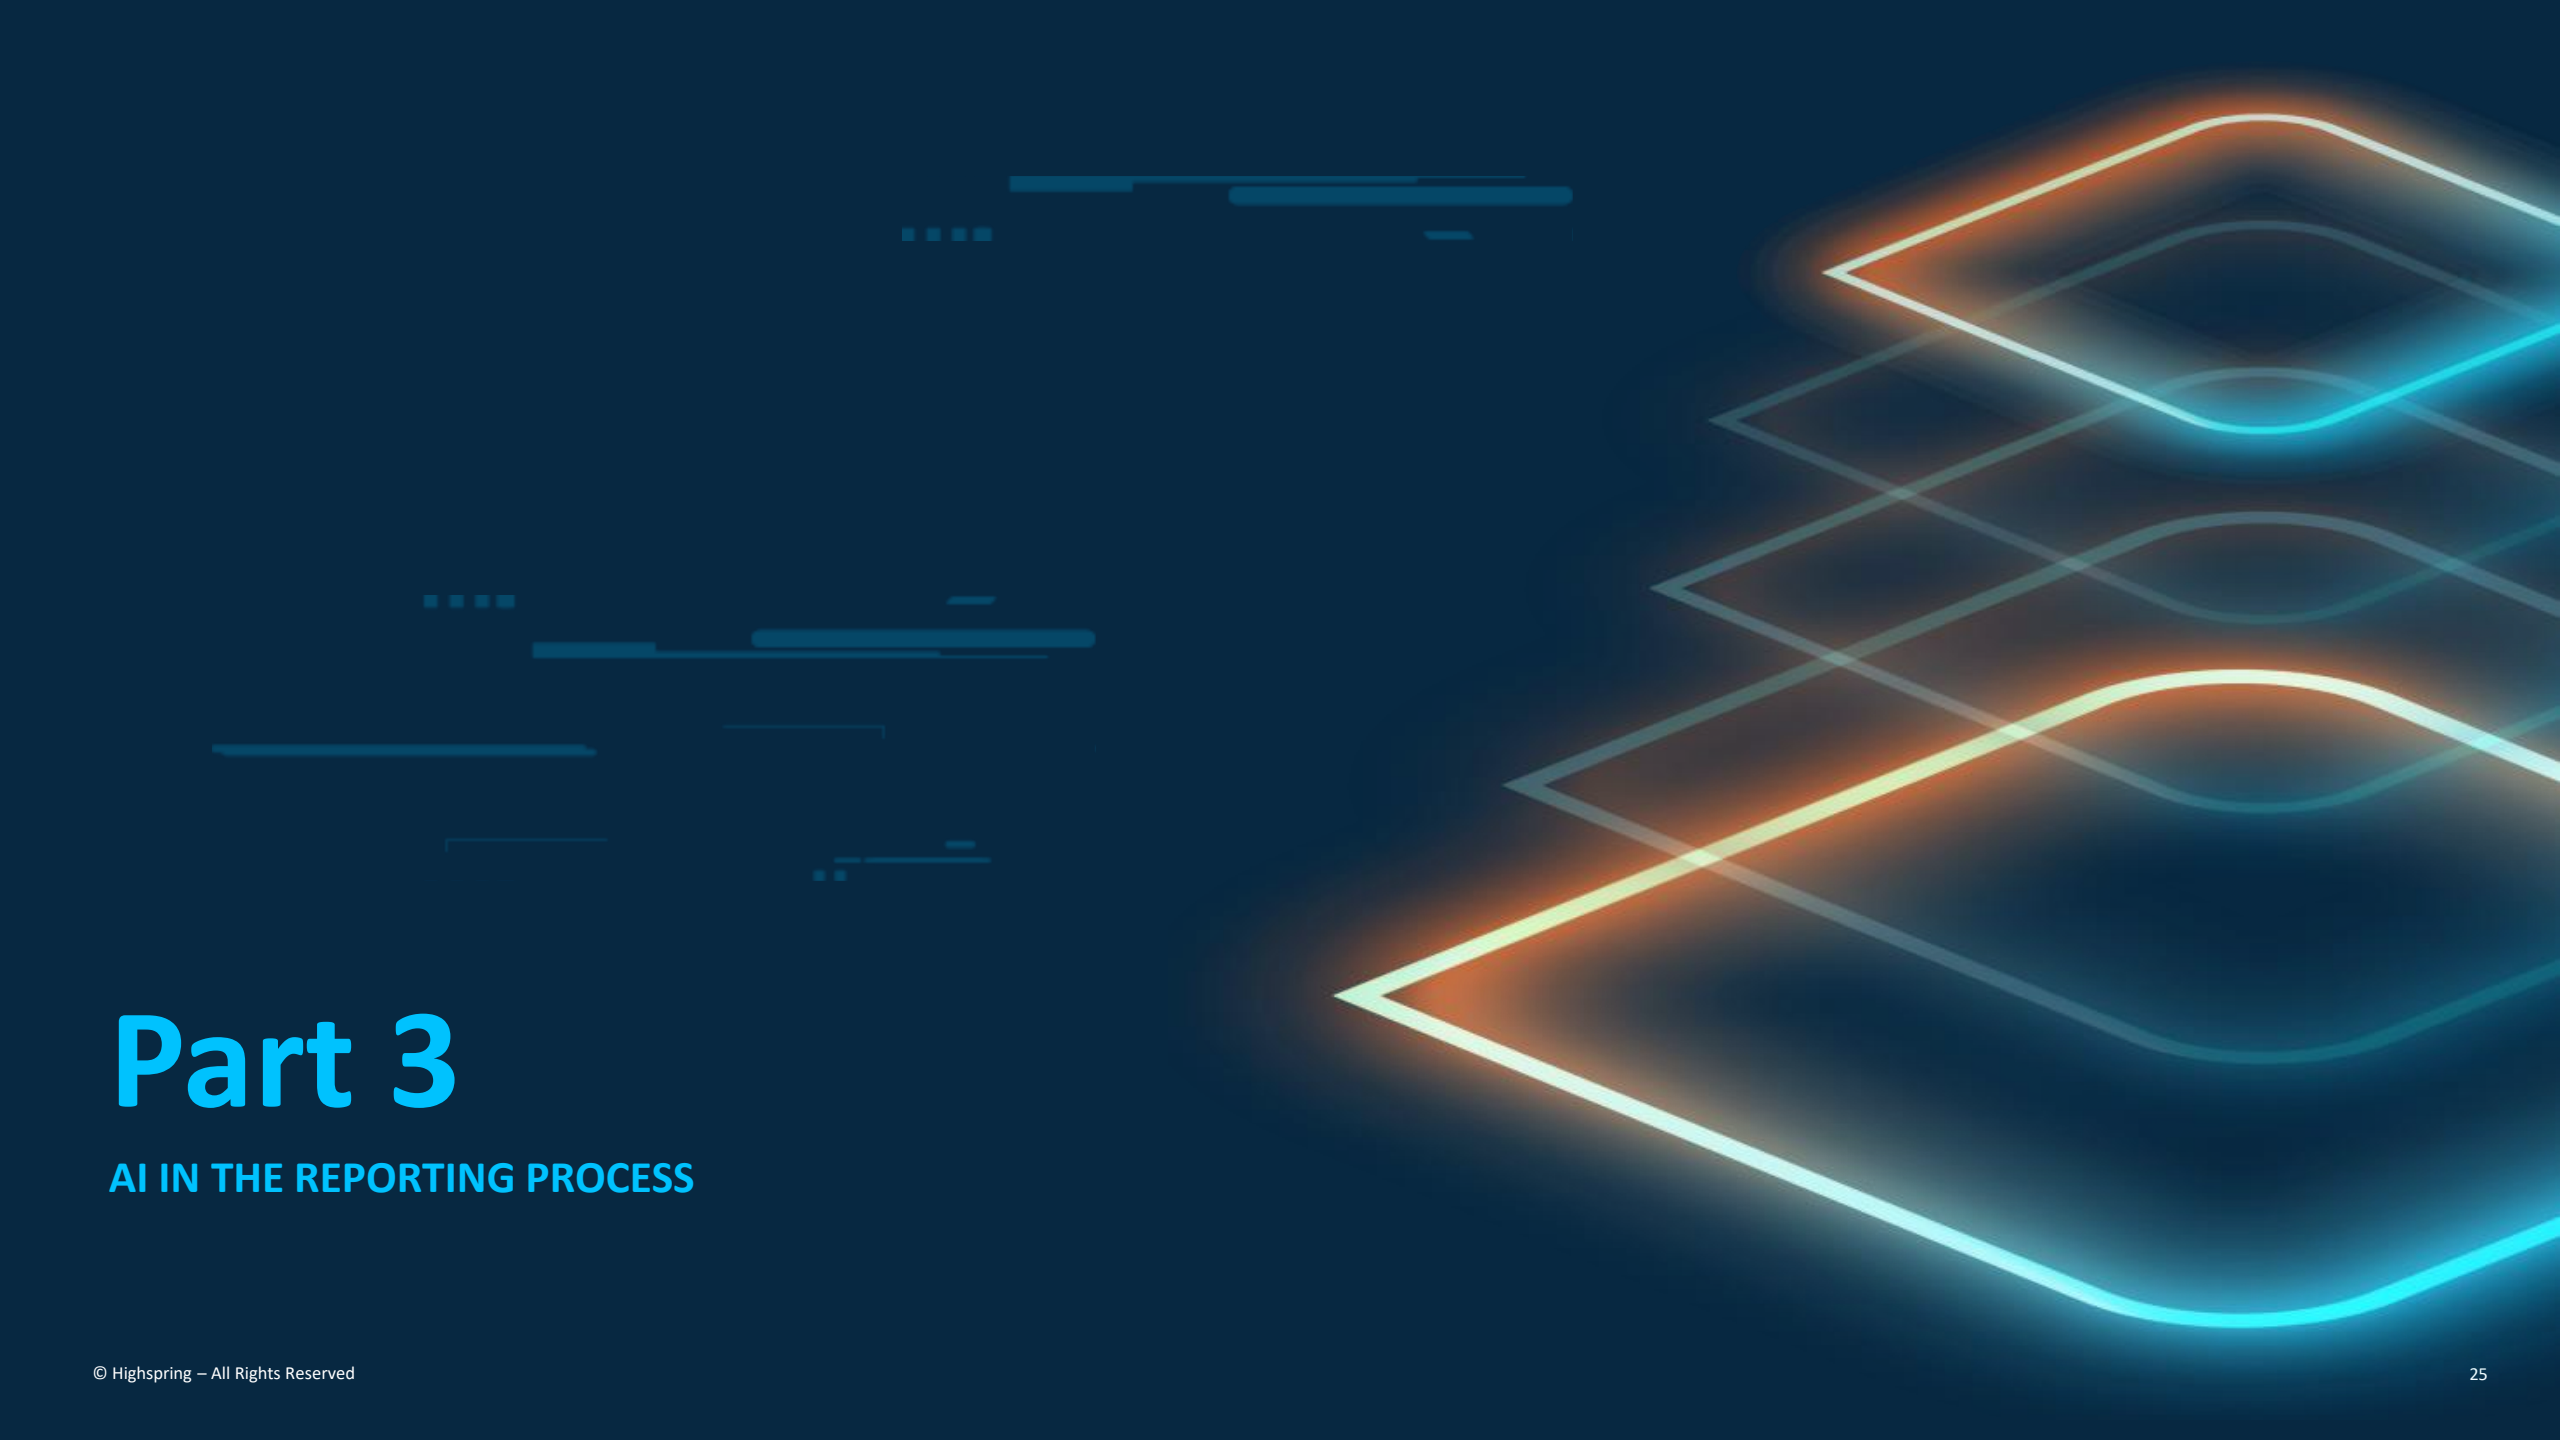

Part 3

AI IN THE REPORTING PROCESS

PART 3: REPORTING CONTROLS

When Does AI Reach the Financials?

For 20 years, controls over reporting systems assumed that once a system is in production it behaves the same way every time. AI doesn't, and that changes how you scope it.

TRADITIONAL SYSTEMS (DETERMINISTIC)

- Same input → same output, every time
- Stability assumed once in production
- ITGCs designed to manage change (access, change mgmt, operations)
- **Risk: unauthorized change**

- Same input → outputs that can vary
- Variation is a feature, not a defect
- Traditional ITGCs are insufficient — output variation also requires control
- ***Risks: drift—change; hallucination—invention; opacity—unclear logic***

The question isn't whether AI is in your environment. It's whether AI's variation reaches a financial assertion without a deterministic check.

THE 3-QUESTION SCOPING TEST

1. **WHERE** does the AI output sit in the process?
   Upstream input? In the workflow? Producing the financial record itself?
2. **WHAT** is the AI producing?
   Data extraction? Classification? Calculation? Judgment?
3. **WHAT** sits between the AI output and the financial record?
   Automated system control? Human review with defined criteria? Nothing?

KEY PRINCIPLE: AI breaks the deterministic assumption that traditional ITGCs were built on.

If AI output reaches a financial statement assertion without a deterministic control between it and the GL → it is in ICFR scope.
Everything else is governance, not ICFR.

Where AI Can Distort the Numbers

AI doesn't introduce new control objectives, it introduces new failure modes within the ones you already run over financial reporting. Entity-level, technology, and process controls extend cleanly to AI when you know what to look for.

FIVE AI RISKS TO THE NUMBERS

Model Drift
Outputs degrade over time as underlying data shifts. A model that worked in Q1 may not work in Q4.

Bias
Embedded patterns produce systematically skewed outputs. Risk concentrates in classifications and exception scoring.

Opacity
Model logic isn't fully explainable or reproducible. Creates audit evidence challenges — show me how you got that answer.

Hallucination
Confident outputs not grounded in source data. Highest risk in narrative generation and summarization.

Data Dependency
AI doesn't fix your data problems. It amplifies them. Garbage in, garbage out, at scale.

EXTENDING YOUR EXISTING FRAMEWORK

Entity-Level
AI governance policy, accountability for AI use cases, training, intake and approval processes.

ITGCs
Data quality and lineage controls, access management for models and training data, change management for model versions and retraining.

Process-Level
Human-in-the-loop checkpoints, output validation against expected ranges, exception monitoring, evidence retention for AI-driven activities.

KEY PRINCIPLE: Don't reinvent ICFR for AI. Extend it.

The framework is not new. The risks are. The work is identifying where AI introduces variation into existing control objectives — and extending ITGCs and process controls to cover it.

Design Principles for AI in Financial Reporting

KEY PRINCIPLE: Not all AI tools are built the same. These are the principles that separate an ICFR-appropriate tool from one that isn't.

Before evaluating where AI fits in your ICFR program, it helps to understand what a well-designed AI tool looks like in an audit and financial reporting context. These four principles — informed by practical implementation experience — define what good looks like.

01

AI STRUCTURES FACTS; HUMANS CONCLUDE ON EFFECTIVENESS

The tool never tells you whether a control is designed or operating effectively. That judgment belongs to a human and cannot be delegated.

02

UNCERTAINTY IS DISCLOSED

The tool shows where it's unsure and where confidence is low. Polished prose that hides ambiguity is a failure mode, not a feature.

03

REFERENTIAL INTEGRITY

Every output ties back to source data the human can verify. No floating claims, no unsourced conclusions.

04

GAP IDENTIFICATION, NOT INVENTION

When reviewing source material, the tool tells you what's missing from the source. Missing context is information, not a problem to solve with invention.

These principles are not aspirational. They are the baseline test for whether an AI tool belongs anywhere near financial reporting.

Intake and Impact Assessment

INTAKE	IMPACT SCREEN	RISK TIER	CONTROL REQUIREMENTS
KEY ACTIVITIES			
▪ Register AI use case through defined intake channel ▪ Identify business owner and accountable decision authority ▪ Define purpose, decision type, and intended users ▪ Confirm data sources and system dependencies ▪ Screen for prohibited / restricted use categories	▪ Execute screening to assess following attributes: ▪ Data sensitivity (e.g. PII, confidential, regulated data) ▪ Financial reporting and ICFR impact ▪ Regulatory and legal exposure ▪ Decision criticality and user reliance ▪ Model autonomy (assistive vs decisioning vs agentic) ▪ Bias and fairness exposure ▪ Third-party model / vendor dependency	▪ Assign Low / Moderate / High impact tier ▪ Map tier to required validation depth ▪ Define human-in-loop requirements ▪ Set approval authority level ▪ Establish monitoring & review cadence ▪ Determine documentation & evidence expectations	▪ Apply control set based on assigned risk tier ▪ Define testing and validation procedures ▪ Require independent risk review for high-impact cases ▪ Document approval and risk acceptance ▪ Establish monitoring metrics and triggers ▪ Set re-assessment and change triggers

KEY PRINCIPLE: Governance is scaled to risk, not one-size-fits-all

AI Monitoring – Keeping Reporting Controls Current

RE-APPROVAL TRIGGERS

- Model Retraining / Version Change
- Third-Party / Vendor AI Model Update.
- New Data Source
- Scope Expansion
- Regulatory Change

MONITORING & CONTROL FOCUS AREAS

- Traditional ITGC Monitoring (access, change management, computer operations)
- Output Anomalies
- Usage Drift
- Bias Indicators
- Exception Patterns
- Human Override Frequency

WHEN TRIGGERS OCCUR – REQUIRED ACTIONS

- Re-assessment routed through AI governance workflow
- Risk / control validation updated as required
- Approval refreshed at appropriate authority level
- Monitoring thresholds & controls recalibrated

TYPICAL REVIEW CADENCE

HIGH IMPACT: QUARTERLY	**MODERATE:** SEMI-ANNUAL	**LOW:** ANNUAL

KEY PRINCIPLE: AI Approval is Conditional — Not Permanent.

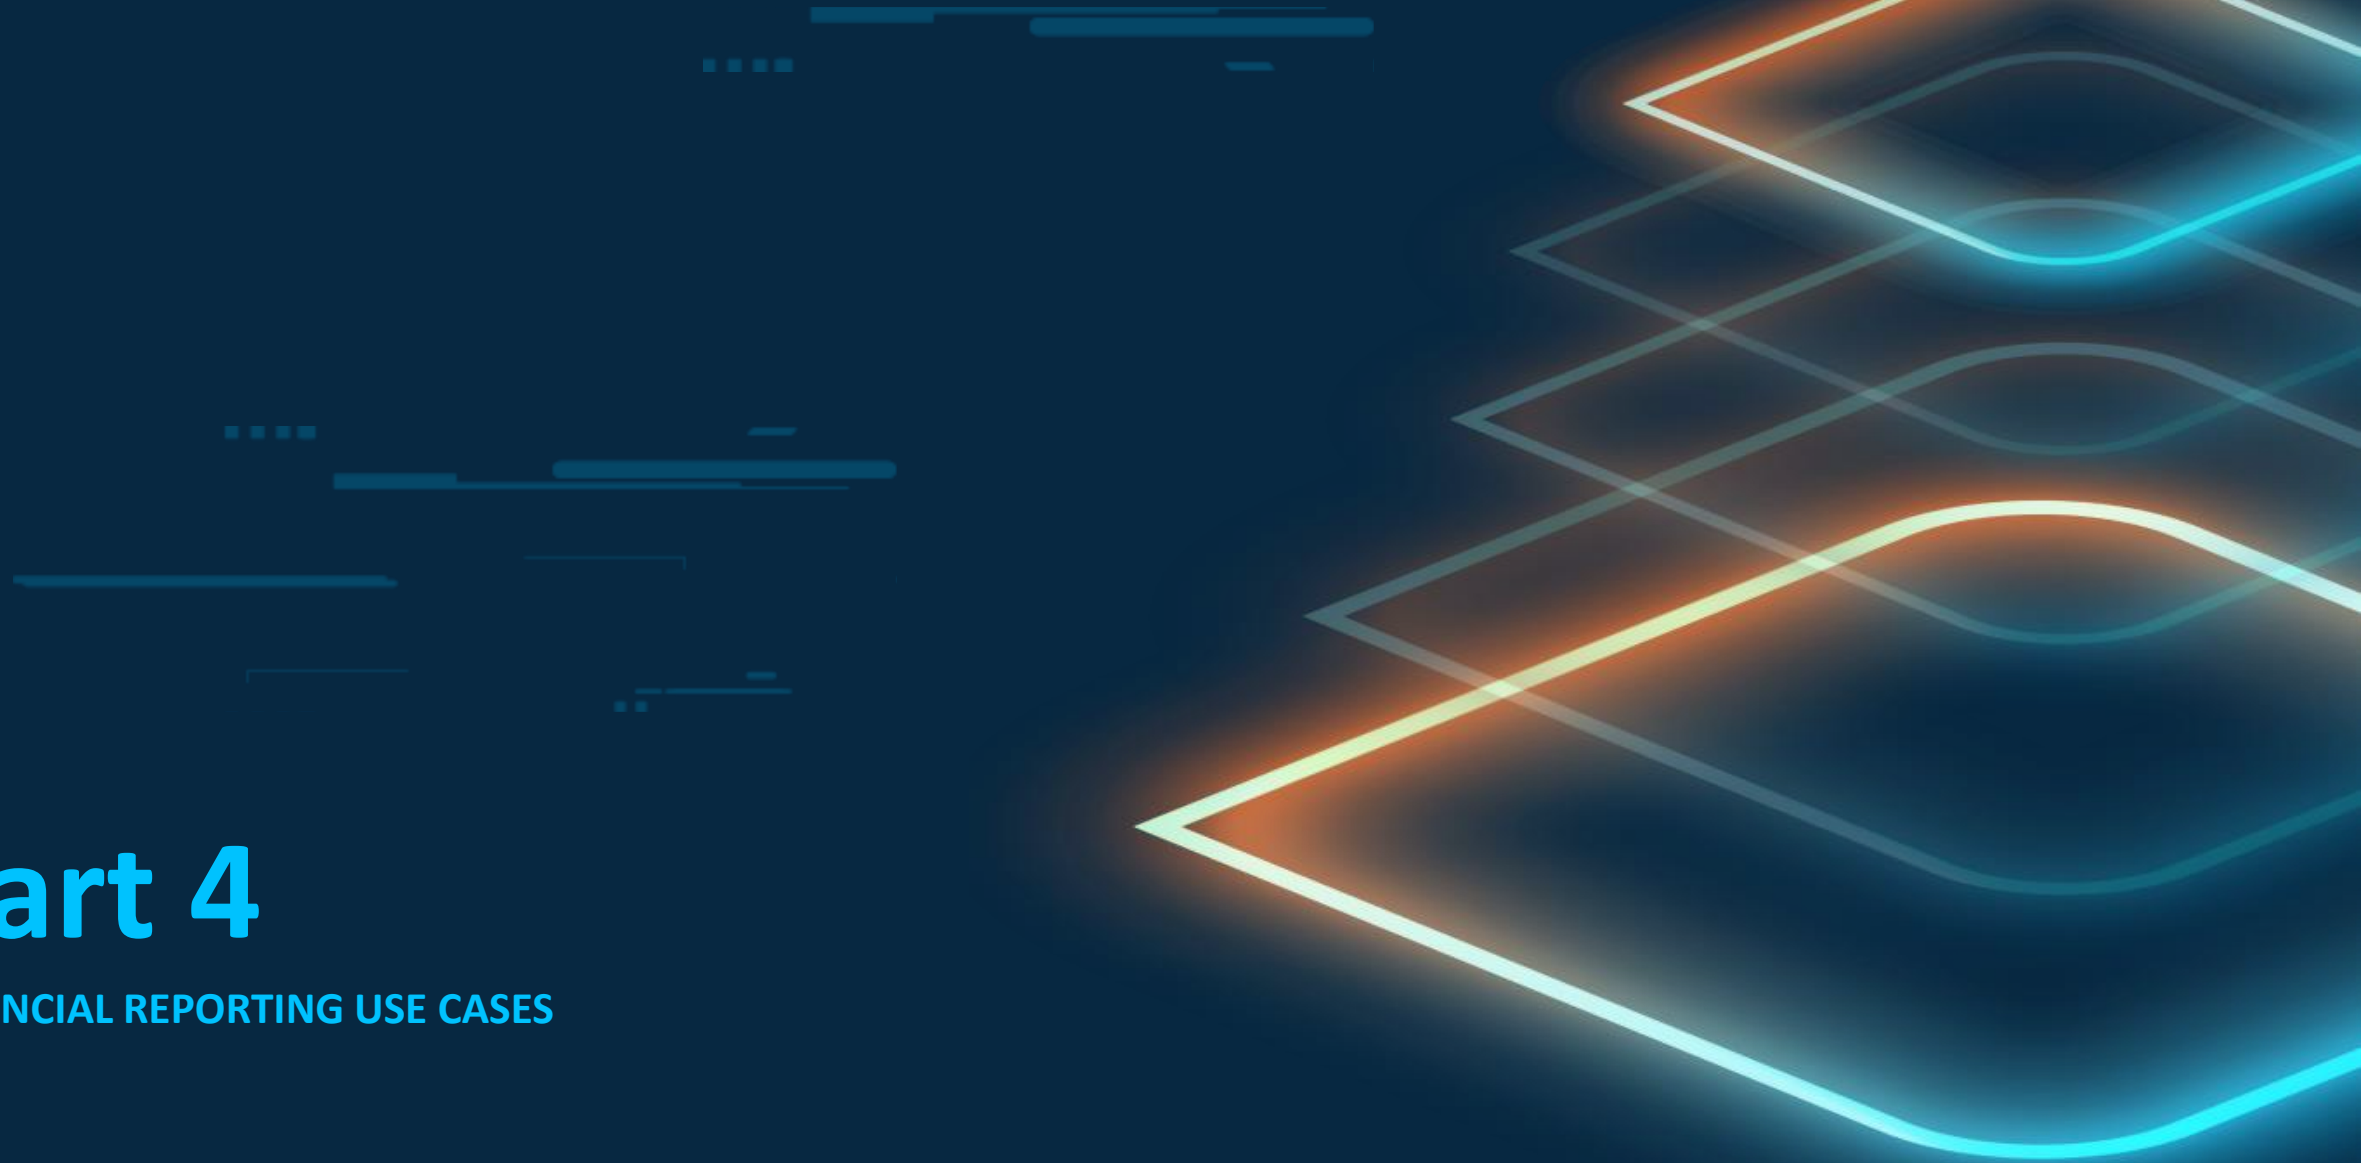

Part 4

FINANCIAL REPORTING USE CASES

PART 4: FINANCIAL REPORTING USE CASES

Where AI Is Showing Up in Financial Reporting

AI handles volume and pattern recognition; humans retain judgment and accountability.

JOURNAL ENTRY & TRANSACTION RISK MONITORING
Full-population JE testing using rules and ML scoring; flags unusual timing, users, accounts, and threshold behavior.

AI-ASSISTED RECONCILIATIONS
ML-based fuzzy matching learns historical clearing patterns; control emphasis shifts to exception review.

INTELLIGENT DOCUMENT PROCESSING FOR IPE
OCR and NLP extract data from invoices and contracts; ties unstructured data to ERP records with confidence scoring.

CLOSE AND FLUX ANALYSIS SUPPORT
AI drafts variance and flux explanations from ledger and sub-ledger data; preparers and reviewers retain the conclusion.

CONTINUOUS CONTROLS MONITORING
Ongoing monitoring of key controls and data feeds between formal testing cycles; flags breakdowns and unusual outcomes.

DISCLOSURE AND NARRATIVE DRAFTING
Drafts footnote and MD&A language from approved figures with source citations; finance retains authorship and review.

Where AI Is NOT Being Used in Financial Reporting

Driven by audit evidence standards and ICFR accountability:

- ✗ Autonomous execution or approval of material transactions
- ✗ Judgment-heavy areas (estimates, reserves, policy decisions) without human review
- ✗ Untethered AI in sensitive reporting workflows (hallucination and data-leakage risk)
- ✗ Black-box models that cannot explain how conclusions were reached
- ✗ Unreviewed AI text in filed disclosures or external reporting

KEY PRINCIPLE: AI is augmenting ICFR — not replacing it.
The wins are in coverage and consistency, not autonomy.

Finance use cases | The art of the possible

AI doesn't redefine what FP&A is responsible for. It expands what's possible. The examples below illustrate common FP&A capabilities that can be enhanced with AI to improve scale, speed, and insight. They represent the art of the possible when AI is applied on top of strong data and process foundations, not a prescriptive set of implementations.

AI-enabled capability	Traditional limitations	How AI can augment	Illustrative impact
Error and fraud detection	**Incorrect financial data or fraudulent transaction** are inputted into an ERP system, introducing financial risks and/or incorrect data reporting into monthly or quarterly board reporting	AI performs a scan on various data sources like payment history and IP addresses to see anomalies. AI can notify a company of fraudulent or incorrect transactions before they are approved	• 80% to 90% reduction in fraudulent or incorrect transactions
Forecasting and predictive analytics	**Many financial models are typically limited to internal data,** constraining the ability for a company to understand significant external factors (e.g., oil prices, geopolitical threats) on future financial performance	Use generative AI/ML modeling, companies can predict forecast trends based on various data inputs including historical trends and current market data to make better business decisions and investments	• Significant improvement in customer churn and sale forecasting • ~ 50% reduction in time vs. traditional analytics processes
Stress-testing and financial scenario modeling	Traditional financial modeling solutions are typically run on a quarterly basis, and **users don't apply a variety of different market scenarios** to understand financial risks (e.g., 2023 high inflation)	AI such as generative adversarial networks can simulate hypothetical, yet plausible, scenarios that are based on complex interdependencies learned from training data without the data limitations of normal financial models	• 20% to 30% improved coverage of various historical financial scenarios • Ability to capture various cross-asset dependence patterns on financial models
Automated cash forecasting and collection management	**Inefficient payment collection strategy negatively impacts cash-flow inefficiency and forecasting.** Continued poor free case flow management can drive potential credit ratings downgrades	AI can improve the efficacy of the forecasting models and collection management by incorporating additional data sources (e.g., CRM, customer emails) to identify problematic customers proactively	• 15 days to collect receivables with uncollectible balances at 0.70%

Automation and AI: Distinct capabilities, stronger together

Automation

- Data integration and validation
- Allocations, roll-forwards, and controls
- Repeatable, rules-based reporting

Automation ensures consistency, speed, and reliability.

Artificial intelligence (AI)

- Pattern and driver identification
- Scenario modeling and sensitivity analysis
- Variance explanation and insight

AI scales analysis and insight, not execution.

Automation + AI

- Clean, validated data feeds advanced analysis
- Automated workflows operationalize insights
- Faster reforecasting with human judgment intact

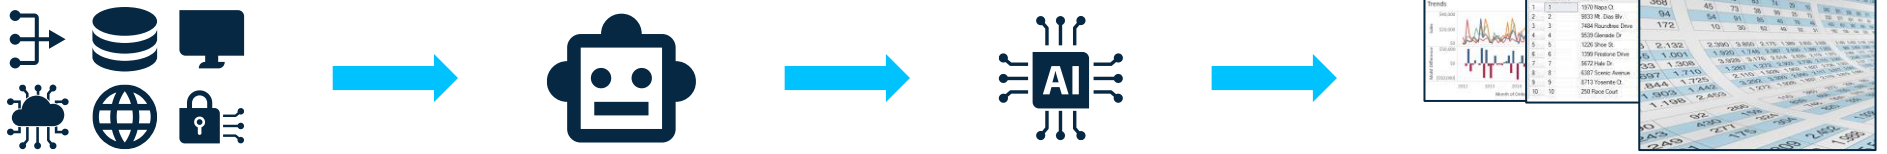

Data sources → Automation → AI → Decision-ready insights

Automation handles execution. AI enhances insight. Finance owns decisions

Financial reporting use cases | Where AI earns its place

Finance owns the numbers and the narrative. These are the places AI is delivering value across the reporting cycle today — and the guardrails that keep the output audit-ready.

AI in the reporting cycle

Where AI does the work

CLOSE & RECONCILIATIONS

Fuzzy matching and auto-clearing across high-volume accounts · exception-based review replaces manual tie-outs.

DISCLOSURE DRAFTING

First-draft footnote and MD&A language built from approved figures with source citations.

VARIANCE & FLUX ANALYSIS

Draft variance explanations from ledger and sub-ledger data · trend analysis across prior periods.

What finance must control

Where the reporting risk sits

OWNERSHIP

A named business owner for each reporting use case · risk tiering and human-in-the-loop confirmed before reliance.

CONTROLS

ITGCs extended to models, source data, and version changes that touch the ledger.

MONITORING

Drift, exception, and override reporting reviewed each close · third-party and vendor AI dependencies assessed.

WHAT MUST STAY TRUE IN EVERY CASE

Practical guardrails whenever AI touches the financial statements

PREPARER JUDGMENT

The controller retains judgment. AI proposes; finance concludes.

TRACEABILITY

Every output traceable to source data and source documents.

EVIDENCE

Evidence retained for the AI-assisted step, not just the final number.

INDEPENDENT REVIEW

A qualified reviewer independent of the AI output signs off before it reaches the ledger.

SCOPE DISCIPLINE

AI use is documented in the ICFR scope and approved before reliance.

The 90-day roadmap | From inventory to controlled scale

Governing AI in financial reporting is less about technology than discipline. Know where AI is already in use, put guardrails around it before it scales, then prove the controls work.

Focus and assess
30 days

- Inventory AI already in use across finance
- Name the business owner for each
- Apply the 3-question ICFR scoping test
- Baseline data and control readiness

Pilot with guardrails
60 days

- Stand up intake and impact screening
- Assign risk tiers and review rules
- Extend ITGCs: data, access, model change
- Set approval authority

Expand and measure
90 days

- Test the AI-related controls
- Monitor drift and exceptions
- Document evidence; scale or stop

Start where reporting risk is highest
- Close and reconciliations
- Journal entry review
- Disclosure drafting

What to Do Next

1. **Inventory where AI is already in your finance function.**
   Even informal use counts. You can't govern what you can't see.

2. **Apply the 3-question scoping test to your top three AI use cases.**
   Identify which ones touch ICFR and which are governance-only.

3. **Extend your existing ITGCs — don't build new ones from scratch.**
   Start with data quality, access, and change management for your AI models.

4. **Establish clear ownership for AI risk in finance.**
   Pick one accountable owner before the second-line and external auditor conversations start.

5. **Engage your external auditor early.**
   They are forming their AI audit approach right now. Better to shape the conversation than react to it.

KEY PRINCIPLE: The smallest useful step beats the perfect plan.

Questions?

Christian Heffron | Partner, Highspring
Email: christian.heffron@highspring.com **linkedin.com/in/christianheffron**

José Ortega Carrero | Partner, Highspring
Email: jose.ortegacarrero@highspring.com

Highspring | vaco BY HIGHSPRING

AGILITY AT WORK.

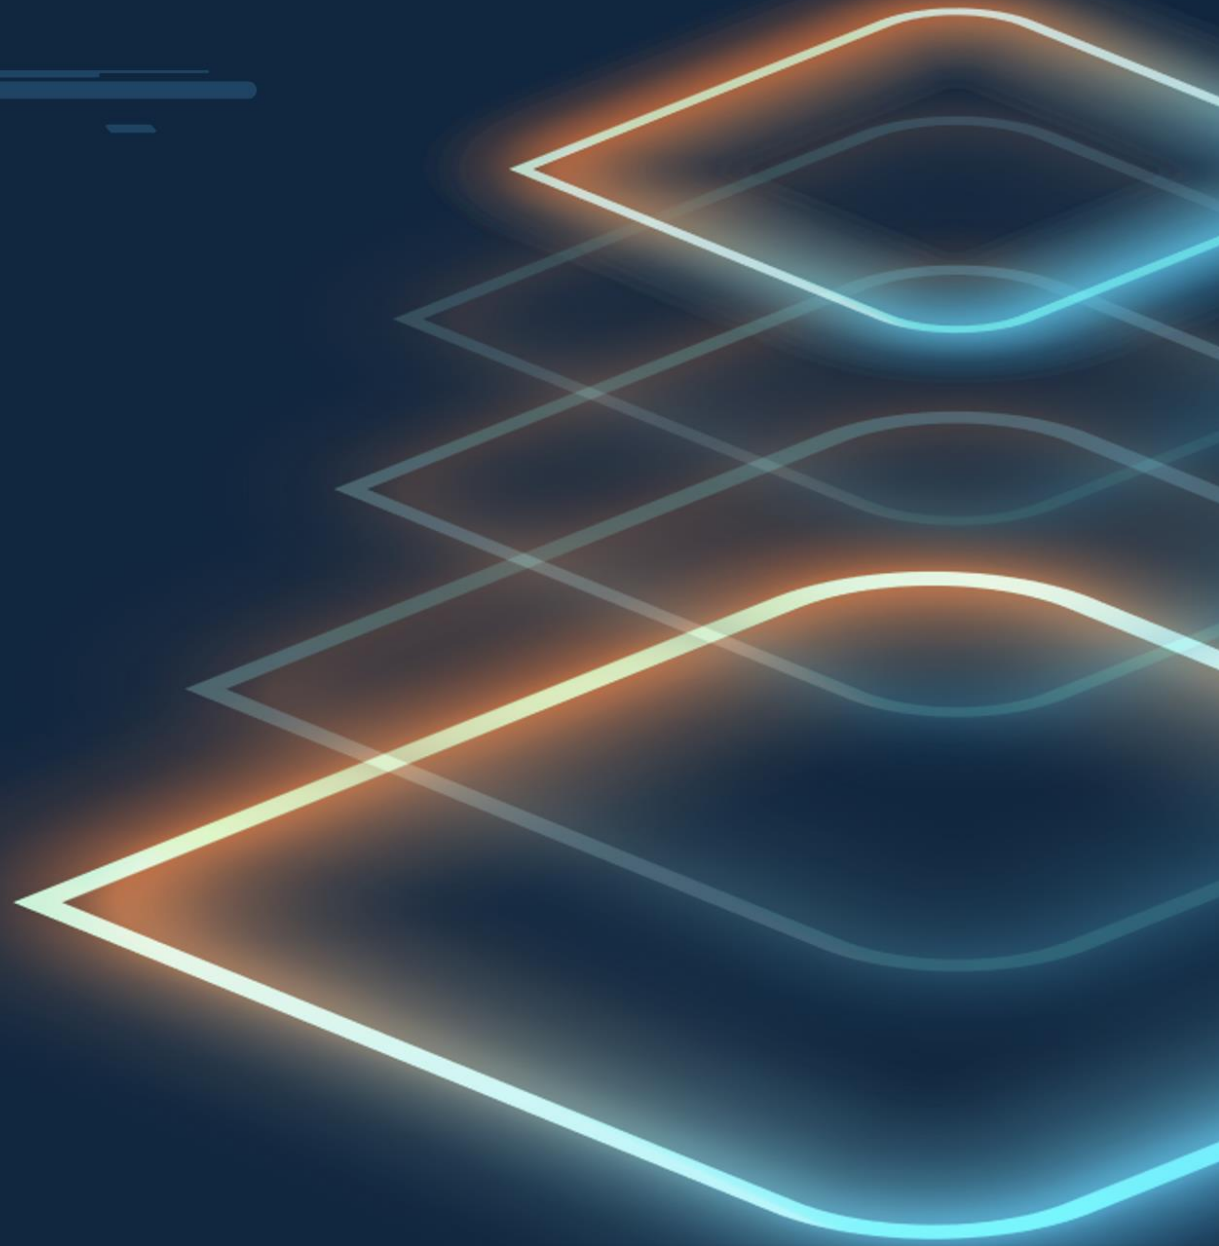

Thank you.