THE FUTURE-READY CFO MASTERCLASS

MODULE 10 OF 20 | PILLAR III: GOVERNANCE, RISK AND RESILIENCE

Enterprise Risk, Cybersecurity and AI Governance

MODULE SUMMARY

This module gives the chief financial officer a practical view of enterprise risk. It covers risk appetite, the risk register, scoring, expected loss and key risk indicators, with a Python verified heat map. It then treats model risk, the risks of artificial intelligence such as hallucination, prompt injection and agentic action, the NIST frameworks, the EU AI Act, privacy, cyber cost, insurance, ransomware, SEC disclosure and vendor dependence. A fictional case, an AI governance operating model with verified inventory scoring, boardroom questions, a risk register extract and an approval checklist follow. A ten question self assessment closes this critical module.

MODULE AT A GLANCE

Core question	How does a CFO see the risks that can damage the enterprise, including cyber, vendor and artificial intelligence risk, decide how much of each to accept, and make sure that controls and evidence stand behind the decision?
Why it matters to the CFO	These risks arrive faster than the annual cycle of a risk committee and land on the income statement and the disclosure record. The CFO owns the numbers that connect risk to capital and insurance.
Key sources	NIST AI RMF and CSF 2.0, SEC cybersecurity rules, the EU AI Act page, COSO, ISO 31000, SR 11-7, IBM (2026), PwC, EY, Deloitte and KPMG.
Executive artifact	An enterprise risk register extract and an AI use case approval checklist, each completed for a fictional company.

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

1. Executive overview and learning objectives

The core question. How does a chief financial officer see the risks that can damage the enterprise, decide how much of each risk the company should accept, and make sure that people, controls and evidence stand behind that decision? The question has always belonged to finance in part, because every risk eventually appears as a loss, a charge, a covenant problem or a disclosure, and because the CFO controls the capital and the insurance that absorb it. What has changed is the speed and the shape of the exposure. A cyber incident can stop billing in an afternoon, a vendor failure can remove a core system for a week, and an artificial intelligence tool can send confidential data outside the company or approve a payment before any person has looked at it.

Why this matters over the next five years. Three developments make the subject urgent. First, the cost of failure is rising. IBM reported in its Cost of a Data Breach Report 2026 that the global average cost of a breach reached 4.99 million dollars, up 12 percent from the prior year and a record, and that organizations that used security artificial intelligence and automation extensively saved 1.93 million dollars compared with those that used none. Second, adoption is moving faster than control. EY reported in August 2025, from a survey of 975 senior executives, that 76 percent of companies were using or planned to use agentic AI, which means software that takes actions on its own, within a year, while only 56 percent were familiar with its risks. Deloitte reported that only about one in five companies had a mature model for governing autonomous agents. Third, the rules are tightening. The Securities and Exchange Commission requires public companies to disclose material cybersecurity incidents within four business days of determining that they are material, and the European Union has adopted a law on artificial intelligence that applies in stages through 2028. These sources describe surveys and rules, they do not measure the losses of any one company, and the reader should treat the statistics as a description of the setting in which the CFO will be asked to approve budgets and sign disclosures.

Terms used in this module. Enterprise risk management is the practice of identifying the events that could prevent the company from reaching its goals, judging how likely and how costly each is, and choosing a response. **Risk appetite** is the amount and type of risk the board is willing to pursue, and **risk tolerance** is the limit beyond which the board expects action. A **risk register** is the list of risks with an owner, a score and a response. A **key risk indicator** is a measure that warns that a risk is rising. **Model risk** is the chance of harm from decisions based on a model that is wrong or misused. A **large language model** is a program trained on text that predicts and generates language, and it is the engine behind generative AI. A **hallucination** is an output that is fluent and false, and a **prompt injection** is an attack in which instructions hidden in the input take control of the model. **Third party risk** is the exposure created by suppliers on whom the company depends, and a **fourth party** is a supplier of that supplier.

What the module covers. Section 2 sets out the concepts: the frameworks of risk management, appetite and tolerance, scoring, the three lines model, the risks specific to artificial intelligence, the NIST frameworks for AI and for cybersecurity, the EU AI Act at a high level, the SEC rules on cybersecurity disclosure and the main privacy regimes. Section 3 works fictional examples that were computed in Python, a programming language widely used for financial modeling: a scored risk register with a heat map and expected loss, a control choice measured against the appetite of the board, key risk indicators, the cost of an incident, a ransomware decision, insurance, model drift, forecast error and fairness tests, the scoring of an inventory of AI use cases, vendor concentration and exit planning. Section 4 places a fictional CFO before the approval of an agent that moves money, Section 5 prepares the boardroom, Section 6 supplies two executive artifacts, and Section 7 reports the research that was read.

How the module connects to the others. Module 1 explains how these AI systems work and where they fail, which this module turns into controls. Module 3 supplies the data discipline on which risk measures depend.

Module 4 evaluates technology spending, including security investment, and this module supplies the risk case for it. Module 9 describes the board and its audit committee, which oversee the risks treated here. Module 11 treats reporting, internal control and assurance, including the extension of assurance to automated processes. Module 12 treats the crisis itself, the playbooks for cash, lenders and communication when a risk becomes an event. Readers who wish to go further should consult the Master Classes in this series on cybersecurity and information risk and on the CFO's AI governance and risk framework.

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

1. Explain the purpose of enterprise risk management, state the difference between risk appetite, tolerance and capacity, and build a scored risk register with owners and key risk indicators.
2. Compute expected loss, read a heat map critically, and test a portfolio of risks and a choice of controls against the appetite set by the board.
3. Describe model risk and the three lines model, and name the AI risks of hallucination, prompt injection, data leakage, intellectual property, bias, over reliance and agentic action, with the controls that answer each.
4. Use the NIST AI Risk Management Framework and the NIST Cybersecurity Framework 2.0 as checklists, and state what the EU AI Act and the SEC cybersecurity rules require at a high level.
5. Judge the cost of a cyber incident, the role and limits of insurance, and the decision framework for a ransomware demand, and know when to bring in counsel, insurers and law enforcement.
6. Manage third party and fourth party concentration, design the AI governance operating model of policy, inventory, approval, monitoring and incident response, and link all of it to continuity and to the board.

2. Conceptual foundations and frameworks

Risk is uncertainty that matters to an objective. A company that sets a goal, such as a margin, a launch date or a covenant, has implicitly accepted every event that could prevent it. Enterprise risk management asks for those events to be written down, owned and compared on one scale, so that the board decides how much to accept and does not discover the answer by accident. Two frameworks dominate. COSO, the Committee of Sponsoring Organizations of the Treadway Commission, published the updated Enterprise Risk Management framework, titled Integrating with Strategy and Performance, in 2017, replacing its original of 2004. ISO 31000, an international standard first published in 2009 and in its second edition since February 2018, gives principles, a framework and a process of identifying, analyzing, evaluating, treating, monitoring and communicating risk, and it is guidance and not a certifiable standard. Two newer documents extend the method to the topics of this module: the NIST Cybersecurity Framework 2.0 and the NIST AI Risk Management Framework. Exhibit 1 sets out how each should be used.

Exhibit 1 Frameworks and Standards Used in Enterprise Risk, Cyber and AI Governance

Framework	What it is	What the CFO takes from it	Limit to remember
COSO Enterprise Risk Management (2017)	A framework that ties risk to strategy and to performance, with guidance on risk appetite, key risk indicators, cloud and cyber risk and artificial intelligence	A common language for the board and a way to connect risk to the planning cycle and to capital (Module 5)	Describes what good looks like and does not supply a method of measurement

Framework	What it is	What the CFO takes from it	Limit to remember
ISO 31000 (2018)	International guidelines on principles, framework and process, usable by an organization of any size and sector	A process that any risk owner can follow: identify, analyze, evaluate, treat, monitor, communicate	Not certifiable, and silent on the specific hazards of AI and of cyber
NIST Cybersecurity Framework 2.0 (2024)	Six functions, Govern, Identify, Protect, Detect, Respond and Recover, for any organization	A checklist for the questions a CFO should put to the CISO, and a basis for reporting to the audit committee	Voluntary, and says what outcomes to reach and not how much to spend
NIST AI Risk Management Framework 1.0 (2023)	Voluntary framework with four functions, Govern, Map, Measure and Manage, for risks of AI to people and organizations	A structure for an AI policy, an inventory and an approval process	Under revision at the time of writing, so confirm the current text
Federal Reserve SR 11-7 (2011)	Supervisory guidance on model risk management for banks, built on validation and effective challenge	The best known statement of how to control a model, widely borrowed outside banking	Written for banks and for statistical models, and not for generative AI as such

Source: COSO, ISO, NIST and Federal Reserve pages opened for this module (Section 7). The column on what a CFO takes from each is the author's view.

Appetite, tolerance and capacity. The three words are often used as if they meant the same thing, and the confusion causes real harm, since a register that has no appetite to compare against is only a list. **Risk capacity** is the loss the company can absorb before its solvency, covenants or franchise are threatened, and it is set by the balance sheet. **Risk appetite** is the board's choice to take a smaller amount than capacity in pursuit of returns. **Risk tolerance** is the specific limit, expressed in a number, at which the board expects management to act. A good appetite statement is short, is stated in measures that finance can compute, and includes at least one limit on the loss from a single event and one on the total of expected losses. Exhibit 2 shows statements for the fictional company used in this module.

Exhibit 2 Capacity, Appetite, Tolerance and Triggers, with Examples for a Fictional Company

Term	Meaning	Example for Harbourline Freight Systems (fictional)	Who sets it
Risk capacity	The loss that the enterprise can bear before solvency, covenants or the franchise are threatened	Covenant headroom permits a one time loss of about 45, which is half of annual EBITDA of 90	Set by the balance sheet, confirmed by the CFO
Risk appetite	The amount of risk the board chooses to carry in pursuit of its goals	Expected annual losses from the registered risks will not exceed 12 percent of EBITDA, which is 10.8	Board, on the proposal of management
Risk tolerance	The limit on a given measure at which management must act	No single event may cost more than 25 percent of EBITDA, which is 22.5, after insurance	Board or its audit committee
Key risk indicator trigger	The level of a leading measure at which a response begins	More than 25 critical vulnerabilities open beyond 30 days is Red and goes to the audit committee	Risk owner, approved by the committee

Fictional. 12 percent of 90 is 10.8, and 25 percent of 90 is 22.5. The thresholds are examples and are not recommended standards; each board must set its own.

Scoring risks on a common scale. To compare a cyber risk with a market risk, each is scored for likelihood and for impact on a scale of one to five, and the two scores are multiplied. The scales in Exhibit 3 are defined by the annual probability of the event and by the loss in millions of dollars if it occurs. The scales should be anchored in money, because words such as "moderate" mean different things to different people, and the anchors for this fictional company are scaled to its earnings of 90 million dollars. A multiplied score has well known weaknesses, which Section 3 demonstrates: many different risks receive the same score, and an event of high impact and low likelihood scores the same as an event of low impact and high likelihood, although the board may feel very differently about them.

Exhibit 3 Likelihood and Impact Scales Anchored in Probability and Dollars

Score	Likelihood in one year	Impact if it occurs (millions of dollars)	Reading of the product
1	Under 5 percent	Under 1	1 to 4: Low, monitor
2	5 percent to under 15 percent	1 to under 3	5 to 9: Moderate, owner manages
3	15 percent to under 35 percent	3 to under 10	10 to 14: High, committee review
4	35 percent to under 60 percent	10 to under 30	15 to 25: Critical, board attention
5	60 percent or more	30 or more	Any event above the tolerance goes to the board

Fictional scales for a company with EBITDA of 90 million dollars, shown as a teaching example. Boundaries are stated as from the lower value up to but not including the upper value.

Model risk and the three lines. Every forecast, credit score, pricing engine and AI assistant is a model, and each can be wrong in a systematic way. The Federal Reserve defined model risk in its guidance SR 11-7, issued in April 2011, as the potential for adverse consequences from decisions based on incorrect or misused model outputs and reports. The guidance rests on **effective challenge**, which it describes as critical analysis by objective, informed parties who can identify the limits of a model and bring about changes, and it names three parts of validation: an evaluation of conceptual soundness, ongoing monitoring, and an analysis of outcomes against actual results. The **three lines model** is a widely used way of assigning the work. The first line, the business and technology teams that build and use a model, owns the risk and the daily controls. The second line, a risk or model validation function, sets the standards and challenges the first. The third line, internal audit, gives the board independent assurance on whether the first two are working. The words **drift** and **bias** name the two failures that matter most in practice. Drift is a change in the data or the world that makes a model less accurate over time, and bias is a pattern of errors that falls more heavily on one group. Exhibit 4 applies the three lines to cyber and AI.

Exhibit 4 The Three Lines Model Applied to Cyber and AI Risk

Line	Who	Role for cyber and AI risk	Evidence the CFO should see
First line	Business owners, technology, security operations, finance process owners	Owns the use case and the control, runs the tests, escalates incidents	Inventory entry, approval record, monitoring results, testing logs
Second line	Risk management, model validation, compliance, privacy, information security governance	Sets policy and standards, validates models independently, challenges ratings and exceptions	Validation reports, challenge log, exception register, aggregate risk reports

Line	Who	Role for cyber and AI risk	Evidence the CFO should see
Third line	Internal audit, external assessors	Tests whether the first and second lines operate as designed, reports to the audit committee	Audit reports, issue tracking to closure, assurance opinions (Module 11)
Board and committees	Audit committee, risk committee, full board	Sets appetite, receives reports, oversees management (Module 9)	Quarterly dashboard, incident reports, review of the appetite statement

Source: author's adaptation of the widely used three lines model and of the validation components of SR 11-7. The mapping to cyber and AI is the author's view.

The risks that are new with artificial intelligence. Most of the risks of generative AI resemble older risks in a new form, but several are new in kind. A model that writes fluent text can state something false with the same confidence as something true. A model that reads documents can be instructed by the documents. A model that is given tools can act. Exhibit 5 names seven risks, with an example from finance and the control that answers it. The common thread is that the controls rest on people and on limits, such as human approval, data restrictions and caps on the amount a system can move, and not on trust in the model.

Exhibit 5 Seven Risks of Artificial Intelligence, with Examples and Controls

Risk	What it is	Example in finance	Principal controls
Hallucination	A confident output that is false or unsupported	A summary of a lease invents a renewal option, or a draft cites a rule that does not exist	Require sources, check numbers against the ledger, human review before use in any filed or external document
Prompt injection	Instructions hidden in input text, such as an email or invoice, that redirect the model	A supplier invoice contains hidden text telling an agent to change the bank account of the payee	Treat all outside content as untrusted, separate instructions from data, limit what the agent can do, require approval for changes to payee data
Data leakage	Confidential information sent to a tool or model provider that may retain or expose it	A staff member pastes draft results or customer lists into a public tool	Approved tools with contract terms on retention and training, data loss prevention, clear rules on data classes
Intellectual property	Use of material to which the company lacks rights, or loss of its own	A model reproduces licensed text, or proprietary code is submitted to an outside service	Vendor terms on ownership and indemnity, counsel review, rules on what may be submitted
Bias	Errors that fall unevenly on groups of people	A credit or hiring tool approves one group at a lower rate for reasons unrelated to risk	Test selection rates by group, document the purpose, involve counsel on the legal test
Over reliance	People accept outputs without proper challenge, so skill and attention decay	Accountants approve suggested codings without reading, and errors pass for months	Sampling of accepted items, measured override rates, rotation of review duties
Agentic action	A system that takes actions, such as sending, paying or changing records, without a person in each step	An agent releases payments or issues credits on its own authority	Narrow permissions, amount and daily limits, human approval above a threshold, logging, a stop switch

Source: author's synthesis from general knowledge of generative AI risk; the framework language of NIST AI 100-1 informs the control column. The examples are illustrative.

The NIST AI Risk Management Framework. NIST released the AI Risk Management Framework 1.0 on January 26, 2023, as a voluntary framework to help organizations manage risks of AI to individuals, organizations and society. The NIST page notes that the framework is being revised as part of the White House AI Action Plan, and that NIST released a generative AI profile, NIST AI 600-1, on July 26, 2024. The framework document describes four functions. **Govern** is a cross cutting function that cultivates a culture of risk management and runs through the other three. **Map** establishes the context in which risks of an AI system are framed. **Measure** uses quantitative, qualitative or mixed tools to analyze, assess, benchmark and monitor risk. **Manage** allocates resources to the risks that were mapped and measured, including plans to respond to, recover from and communicate about incidents. The document also lists characteristics of trustworthy AI: valid and reliable, safe, secure and resilient, accountable and transparent, explainable and interpretable, privacy enhanced, and fair with harmful bias managed, and it observes that these often involve trade offs that must be balanced in context. Exhibit 6 restates the functions as questions for a finance organization.

Exhibit 6 The Four Functions of the NIST AI Risk Management Framework, Restated for Finance

Function	What NIST describes	Question for the CFO	Evidence in the finance organization
Govern	Culture, accountability, policies and oversight, running through all functions	Who is accountable for each AI system, and what policy governs its use?	AI policy, named owners, committee charter, training record (Module 1)
Map	Context for each system, its purpose, users, data and possible harms	What does each tool do, with what data, and who is affected if it errs?	Inventory of use cases with data classes and impact ratings (Exhibit 22)
Measure	Tools to analyze, assess, benchmark and monitor risk and impact	How do we know it works, and how do we know it still works?	Test results, drift and error measures, bias tests (Exhibit 19 to Exhibit 21)
Manage	Allocation of resources to mapped and measured risks, with response and recovery plans	What do we do when it fails, and who can stop it?	Controls by tier, incident procedure, stop switch, vendor exit plan

Source: NIST AI RMF page and NIST AI 100-1 (Sections 3 and 5), opened for this module. The questions and evidence columns are the author's translation for finance.

The NIST Cybersecurity Framework 2.0. NIST published version 2.0 on February 26, 2024. Its most visible change was the addition of a sixth function, **Govern**, which treats cybersecurity as a major enterprise risk alongside finance and reputation, and which brings the CFO and the board squarely into the framework. The framework also moved beyond critical infrastructure to organizations of any sector. The document defines the functions in this way: Govern sets, communicates and monitors the risk strategy, expectations and policy; Identify builds an understanding of current risks, including those tied to assets and suppliers; Protect applies safeguards such as access control, training and data security; Detect finds and analyzes possible attacks; Respond acts on detected incidents; Recover restores affected assets and operations. Exhibit 7 gives each function a question that a CFO can ask.

Exhibit 7 The Six Functions of the NIST Cybersecurity Framework 2.0, with Questions for the CFO

Function	Purpose as NIST describes it	Question the CFO should ask	A sign of strength
Govern	Sets, communicates and monitors the cybersecurity risk strategy, expectations and policy	Who owns cyber risk at executive level, and does the board see it quarterly?	A stated appetite, a funded plan, and a named executive who reports to the audit committee

Function	Purpose as NIST describes it	Question the CFO should ask	A sign of strength
Identify	Understands current risks, including those tied to assets, suppliers and improvement opportunities	Do we know our critical assets, data and suppliers, and which are most exposed?	A current asset and vendor inventory ranked by criticality
Protect	Applies safeguards such as access control, training, data and platform security	How are privileged access, patching and backups controlled and tested?	Multifactor authentication on privileged accounts, patch targets met, tested backups
Detect	Finds and analyzes possible attacks and compromises in a timely way	How quickly would we know, and who would be told?	Monitoring, a measured time to detect and an escalation path
Respond	Manages, analyzes, mitigates, reports and communicates on incidents	Who decides on payment, disclosure and customer notice, and have they rehearsed it?	A tested playbook with counsel, insurer and communications roles named (Module 12)
Recover	Restores affected assets and operations and supports communication during recovery	How long to restore the processes that produce cash, and have we shown it?	Demonstrated recovery times that meet stated targets (Exhibit 28)

Source: NIST CSF 2.0 (NIST CSWP 29, February 26, 2024) and the NIST announcement, opened for this module. The questions and signs of strength are the author's.

The EU AI Act at a high level. The European Commission page on the regulatory framework states that the AI Act entered into force on August 1, 2024, and sorts systems into four tiers of risk: unacceptable risk, which is banned; high risk, which is regulated; limited or transparency risk, which carries disclosure duties such as telling people that they are dealing with a machine; and minimal risk, which carries no new rules. The page also reports that a digital omnibus on AI, which entered into force on July 27, 2026, delayed the rules for high risk systems. Exhibit 8 lists the dates as the page gives them. A company that sells to, or operates in, the European Union, or that uses systems whose output is used there, should ask counsel whether the Act applies, and the dates and obligations should be confirmed on the current official text, since they have already changed once.

Exhibit 8 EU AI Act, Dates and Categories as Stated on the European Commission Page

Topic	What the European Commission page states	Implication to raise with counsel
Entry into force	August 1, 2024	Obligations apply in stages and not all at once
Prohibited practices	Eight prohibitions applied from February 2025, and a ninth, on non consensual sexually explicit content, takes effect in December 2026 through the omnibus	Confirm that no use case in the inventory falls in a prohibited category
General purpose AI models	Obligations applied from August 2, 2025, with enforcement by the AI Office	Mainly relevant to providers of models, and to the vendor terms of those who use them
Transparency (limited risk)	Rules apply from August 2026, such as telling people they deal with a machine and labelling deepfakes	Applies to customer facing chat and content, so involve legal and marketing
High risk, listed uses (Annex III)	Apply from December 2, 2027, for areas such as employment, education, biometrics and critical infrastructure	Hiring and screening tools deserve early attention (use case U9 in this module)

Topic	What the European Commission page states	Implication to raise with counsel
High risk, regulated products (Annex I)	Apply from August 2, 2028	Relevant to manufacturers whose products embed AI
Penalties	The page does not state amounts	Do not rely on figures from memory, confirm with counsel and the official text

Source: European Commission, digital-strategy page on the AI regulatory framework, opened for this module (Section 7). Nothing beyond the page has been added. This is not legal advice.

The SEC rules on cybersecurity disclosure. The Securities and Exchange Commission adopted rules on July 26, 2023, and its press release states two requirements that bear directly on the CFO. First, a public company must disclose a material cybersecurity incident on Form 8-K under Item 1.05, generally within four business days after it determines that the incident is material. The deadline can be delayed only if the United States Attorney General finds that immediate disclosure would threaten national security or public safety and notifies the Commission in writing. Second, a company must describe in its annual report on Form 10-K, under Item 106 of Regulation S-K, its processes for managing cybersecurity risk, the material effects of risks and incidents, the oversight of the board, and the role and expertise of management. Foreign private issuers must make comparable disclosures on Forms 6-K and 20-F. The practical consequence for the CFO is that the decision about materiality must be made without unreasonable delay, by a defined process that involves legal, security and finance, and that the facts for the disclosure must be gathered at the speed of the incident. Exhibit 9 summarizes the rules.

Exhibit 9 SEC Cybersecurity Disclosure Rules for Public Companies, as Stated in the SEC Press Release

Requirement	What the SEC release states	What the CFO should have ready
Form 8-K, Item 1.05	Disclosure of a material cybersecurity incident, generally due four business days after the company determines the incident is material	A materiality procedure with named decision makers, a template, and a record of the analysis (Module 11)
Delay for national security	Available only if the Attorney General finds that immediate disclosure would threaten national security or public safety and notifies the Commission in writing	Counsel who know the process, since this is not a discretion of the company
Form 10-K, Item 106 of Regulation S-K	Annual description of risk management processes, material effects of risks and incidents, board oversight, and the role and expertise of management	A consistent account of governance that matches what the board actually does (Module 9)
Foreign private issuers	Comparable disclosure on Form 6-K for incidents and Form 20-F for risk management, strategy and governance	Awareness for groups with foreign parents or listings
Smaller reporting companies	Received an additional 180 days before Form 8-K disclosures began	No longer a source of relief for most, since the dates have passed

Source: SEC Press Release 2023-139 (July 26, 2023), opened for this module. Later staff guidance or interpretations were not read and are not described. Public companies should rely on counsel for current practice.

Privacy and data protection. Privacy laws regulate the collection, use, sharing and protection of information about individuals, and they differ by place. The General Data Protection Regulation of the European Union and the California Consumer Privacy Act, as amended, are the best known, and a growing number of other states and countries have comparable laws. The subject is stated here from the author's own knowledge and not from a source opened for this module, and it is limited to the points that a CFO needs in order to ask good questions. The obligations, thresholds and penalties of each law are matters for counsel, and no figure is

given here. Exhibit 10 lists the questions.

Exhibit 10 Privacy and Data Protection, Questions for the CFO to Put to Counsel

Topic	What it means for the company	Question for management and counsel
Scope	Laws often apply according to where the individual lives and not only where the company is located	Which laws reach our customers, employees and website visitors?
Data inventory	A company cannot protect or delete what it cannot find	Do we have a record of what personal data we hold, where, and why?
Purpose and minimization	Data collected for one purpose should not be used freely for another, including for training a model	Do our AI use cases use personal data, and is that use permitted?
Rights of individuals	Individuals may have rights to see, correct or delete their data	Can we find and act on a request within the required time?
Vendors and transfers	The company usually remains accountable for data held by its processors, and cross border transfers may be restricted	Do contracts bind our vendors, and do we know where the data goes?
Breach notification	Many laws require notice to regulators and individuals within set periods after a breach	Is notice built into the incident playbook, with counsel leading?

Source: author's own knowledge, stated at a high level and not confirmed in a source read for this module. Consult qualified privacy counsel for the laws that apply.

3. Practical applications and tools

A. The scored register, the heat map and expected loss. The method has five steps. List the risks in the language of the business, so that each one describes an event and not a topic. Assign an owner who has the authority to act. Estimate the annual probability and the loss if the event occurs, using ranges and evidence where it exists. Convert the estimates to scores with the scales of Exhibit 3, and compute the expected loss, which is probability multiplied by impact. Then compare the total with the appetite. Exhibit 11 shows twelve risks for Harbourline Freight Systems, a fictional company with revenue of 600 million dollars and earnings before interest, taxes, depreciation and amortization (EBITDA) of 90 million dollars. All figures are in millions of dollars and were computed in Python.

Exhibit 11 Risk Register of the Fictional Harbourline Freight Systems (Millions of Dollars)

No.	Risk event	Owner	Prob.	Impact	L	I	L x I	Rating	Expected loss
R1	Ransomware halts dispatch and billing systems	CIO	10%	18.0	2	4	8	Moderate	1.80
R2	Business email compromise redirects a supplier payment	Controller	32%	2.4	3	2	6	Moderate	0.77
R3	Cloud provider outage longer than three days	CIO	8%	12.0	2	4	8	Moderate	0.96
R4	Personal data breach with regulatory action	General counsel	6%	9.0	2	3	6	Moderate	0.54
R5	Error in the AI route pricing model	Chief commercial officer	22%	4.5	3	3	9	Moderate	0.99

No.	Risk event	Owner	Prob.	Impact	L	I	L x I	Rating	Expected loss
R6	Confidential data entered into a public generative AI tool	CISO	45%	2.5	4	2	8	Moderate	1.12
R7	Loss of the largest customer, 24 percent of revenue	CEO	7%	35.0	2	5	10	High	2.45
R8	Fuel and tariff cost shock	CFO	55%	7.0	4	3	12	High	3.85
R9	Failure of a key subcontracted carrier	Chief operating officer	18%	6.0	3	3	9	Moderate	1.08
R10	Loss of key technology leaders	CHRO	12%	1.8	2	2	4	Low	0.22
R11	Breach at a subcontractor of our software vendor (a fourth party)	CISO	14%	5.0	2	3	6	Moderate	0.70
R12	Change in rules on driver classification	General counsel	40%	1.4	4	2	8	Moderate	0.56
	Total expected annual loss								**15.04**

Fictional. Expected loss is probability times impact, for example 0.10 times 18.0 is 1.80 and 0.55 times 7.0 is 3.85. The twelve expected losses sum to 15.04 (15.039 before rounding). Scores from the scales of the preceding exhibit. Counts by rating: Low 1, Moderate 9, High 2, Critical 0.

Reading the register. The scores alone suggest a comfortable picture, since nine of the twelve risks are rated Moderate, two High and none Critical. The expected losses tell a different story. The total is 15.0 million, which is 16.7 percent of EBITDA, and the two risks rated High, the fuel shock and the loss of the largest customer, account for 41.9 percent of it. Ransomware ranks 3 by expected loss but is tied for fifth by score, because three other risks also score 8, and the business email fraud ranks 8 by expected loss and ninth by score, tied with two others at 6. A multiplied score cannot tell a frequent small loss from a rare large one, and ranking by score alone would send management attention to the wrong places. Risks connected to information technology, vendors, privacy and AI account for 7.96 million, or 52.9 percent of the total, which is a reason for the CFO to treat the subject of this module as a principal financial risk and not as a technical matter.

The heat map. A heat map places each risk in a grid of likelihood against impact so that the board can see the whole at a glance. Exhibit 12 is built from the same scores, and its value lies in the conversation it provokes, such as why the ransomware and the cloud outage sit in the same cell, and why the single loss of the largest customer lies in a column of its own. Its weakness is that it shows ordinal ranks and not dollars, and so it should always be accompanied by the expected losses and by the single event limit of the tolerance statement.

Exhibit 12 Heat Map of the Twelve Risks, by Likelihood Score and Impact Score

Likelihood \ Impact	1 Minor	2 Limited	3 Significant	4 Major	5 Severe
5 Almost certain	None	None	None	None	None
4 Likely	None	R6, R12	R8	None	None
3 Possible	None	R2	R5, R9	None	None

Likelihood \ Impact	1 Minor	2 Limited	3 Significant	4 Major	5 Severe
2 Unlikely	None	R10	R4, R11	R1, R3	R7
1 Rare	None	None	None	None	None

Fictional. Cells list the register numbers. Cell counts: likelihood 4 row, R6 and R12 in impact 2 and R8 in impact 3; likelihood 3 row, R2 in impact 2 and R5 and R9 in impact 3; likelihood 2 row, R10 in impact 2, R4 and R11 in impact 3, R1 and R3 in impact 4 and R7 in impact 5. Twelve risks in all.

B. Choosing controls against the appetite. The appetite statement of Exhibit 2 limits expected annual losses to 10.8 million, and the register totals 15.0 million, an excess of 4.2 million. Management proposes five responses, each priced and each reducing either the probability or the impact of one risk. Exhibit 13 computes the expected loss before and after, and the annual cost, and the difference between the reduction and the cost, which is the net benefit measured on expected loss alone.

Exhibit 13 Control Options, Cost and Reduction in Expected Loss (Millions of Dollars)

Control	Risk	Cost	Before: prob. x impact	After: prob. x impact	Loss before	Loss after	Cut	Net
O1 Immutable backups and network segmentation	R1	0.90	10% x 18.0	5% x 12.0	1.800	0.600	1.200	0.300
O2 Call back verification of changed bank details	R2	0.15	32% x 2.4	8% x 2.4	0.768	0.192	0.576	0.426
O3 Approved generative AI tool with data loss prevention	R6	0.35	45% x 2.5	15% x 2.5	1.125	0.375	0.750	0.400
O4 Second region and tested failover for core systems	R3	0.80	8% x 12.0	8% x 4.0	0.960	0.320	0.640	(0.160)
O5 Fuel surcharge clauses and a partial hedge	R8	0.60	55% x 7.0	55% x 4.5	3.850	2.475	1.375	0.775
All five controls		**2.80**					**4.541**	**1.741**

Fictional. Loss is expected loss, Cut is the reduction in expected loss, and Net is the cut less the annual cost. Computed in Python. Example O1: 0.10 x 18.0 = 1.800 falls to 0.05 x 12.0 = 0.600, a reduction of 1.200 against a cost of 0.90. Total reduction 1.200 + 0.576 + 0.750 + 0.640 + 1.375 = 4.541, total cost 2.80.

The result and the test of the appetite. Four of the five controls pay for themselves on expected loss alone, and the second region and failover, O4, does not: it costs 0.80 million and removes 0.64 million of expected loss, a net cost of 0.16 million. The CFO could stop there, but without O4 the total expected loss is 11.138 million, which is still above the limit of 10.8 million. With all five the total falls to 10.498 million, which is 11.7 percent of EBITDA and inside the appetite. Exhibit 14 sets out the test, including the tolerance for a single event.

Exhibit 14 Appetite and Tolerance Test, Before and After Controls (Millions of Dollars)

Test	Limit	Before controls	After all five controls	Result after
Total expected annual loss from registered risks	10.8 (12 percent of EBITDA)	15.04 (16.7 percent)	10.50 (11.7 percent)	Within appetite

Test	Limit	Before controls	After all five controls	Result after
Total without the second region (O4)	10.8	15.04 before, 11.14 after four controls	11.14	Above appetite by 0.34
Largest single event, before insurance	22.5 (25 percent of EBITDA)	R7 loss of largest customer at 35.0	Unchanged at 35.0	Above tolerance, board decision needed
Annual cost of the five controls	Not a limit; compare with reduction	None	2.80 against reduction of 4.54	Net benefit of 1.74 in expected loss

Fictional. Computed in Python: 15.039 less 4.541 is 10.498, which is below 10.8; 15.039 less 3.901 (the four controls other than O4) is 11.138, which is above 10.8. The excess before controls is 15.039 less 10.8, or 4.239.

What the result teaches. Two lessons deserve emphasis. The first is that appetite is a policy choice and not a computation: the board might instead raise the limit, but it should do so knowingly, and the test makes the cost of the choice visible. The second is that the largest single exposure, the customer concentration of R7, is not touched by any control in the list, and it exceeds the tolerance of 22.5 million. It is a strategic risk that belongs in the discussions of Module 6 and Module 8, and the register has done its job by showing that the technology controls, however valuable, do not address the largest item. Expected loss also understates the danger of rare and severe events, since an average of many years hides the year in which the loss occurs, which is why the tolerance on a single event is stated separately.

C. Key risk indicators. A register is reviewed quarterly, and risks move faster than that. A key risk indicator is a leading measure that is cheap to collect and that moves before the loss does. Each indicator should have a green limit and an amber limit set by the owner and approved by the committee, an owner, a source and a trend. Exhibit 15 shows eight indicators for the fictional company, with the status computed by rule. The three Red items are the most useful content on the page, since each is a management failure and not a technical one: vendors without a current assessment, AI use cases running without approval, and staff using unsanctioned tools.

Exhibit 15 Key Risk Indicators for Cyber, Vendor and AI Risk (Fictional)

Indicator	Owner	Unit	Current	Green limit	Amber limit	Status
Critical vulnerabilities unpatched beyond 30 days	CISO	count	14	at most 10	at most 25	Amber
Phishing simulation click rate	CISO	percent	4.1	at most 5	at most 8	Green
Privileged accounts protected by multifactor authentication	CIO	percent	97	at least 99	at least 95	Amber
Restore tests of critical backups that met the target	CIO	percent	92	at least 98	at least 90	Amber
Critical vendors with a current security assessment	Procurement lead	percent	71	at least 90	at least 75	Red
AI use cases in production without an approval record	AI governance lead	count	3	at most 0	at most 1	Red
Mean time to detect a confirmed incident	CISO	hours	19	at most 24	at most 72	Green

Indicator	Owner	Unit	Current	Green limit	Amber limit	Status
Network traffic to unsanctioned AI tools, share of all AI traffic	CISO	percent	11	at most 5	at most 10	Red

Fictional. Status rule: Green if the current value meets the green limit, Amber if it meets the amber limit, otherwise Red. Counts: Green 2, Amber 3, Red 3.

D. The cost of a cyber event, bottom up. The headline averages from industry studies are useful context but cannot replace an estimate for the company. IBM reported a global average cost of 4.99 million dollars per breach in 2026, and the page does not give the composition used here. A bottom up estimate lists the components of cost for a specific scenario, and it should include the items that are usually missed, such as contract penalties and the revenue lost from customers who leave afterward. Exhibit 16 builds the 18 million dollar impact of the ransomware risk R1, which is 3.61 times the global average, a ratio that is plausible for a freight company whose revenue depends on systems that run every day of the week.

Exhibit 16 Bottom Up Cost of the Ransomware Scenario R1 (Fictional)

Component of cost	Millions of dollars	Share of total
Lost contribution during nine days of outage at 0.8 per day	7.2	40.0 percent
Forensics, restoration and rebuilt systems	2.4	13.3 percent
Legal advice, notification and credit monitoring	1.1	6.1 percent
Service credits and contract penalties paid to customers	3.0	16.7 percent
Regulatory response and fines	0.6	3.3 percent
Customer churn in the following year, estimated contribution	3.7	20.6 percent
Total impact assumed for risk R1	**18.0**	**100.0 percent**

Fictional. Computed in Python: 7.2 + 2.4 + 1.1 + 3.0 + 0.6 + 3.7 = 18.0, and 18.0 / 4.99 = 3.61. The global average of 4.99 million dollars is from the IBM Cost of a Data Breach Report 2026 page, opened for this module.

E. The ransomware decision. When systems are encrypted and a ransom is demanded, the CFO is one of a small group who must decide whether to pay. The decision belongs to the board and the chief executive, with the advice of counsel, the insurer, the forensic firm and law enforcement, and nothing in this module is a recommendation to pay or not to pay. Payment may be unlawful if the recipient is subject to sanctions, it does not guarantee a working decryption key, and it may mark the company as willing to pay. The financial framework that supports the discussion is simple, and it compares the expected cost of each path in downtime and ransom, on the assumptions that must be tested in advance. In the fictional case the demand is 4.0 million dollars, each day of outage costs 0.8 million, a decryption key works with a probability of 0.6 and then needs 5 days to restore, and if it fails the company loses 3 days before it falls back on its backups. Exhibit 17 shows how the answer depends on the time needed to restore from backups.

Exhibit 17 Expected Cost of Paying and Not Paying a Ransom, by Backup Restore Time (Millions of Dollars)

Days to restore from backups	Cost of not paying	Cost of paying, expected	Paying less not paying	Cheaper on this measure
6	4.80	9.28	4.48	Not paying
9	7.20	10.24	3.04	Not paying
14	11.20	11.84	0.64	Not paying

Days to restore from backups	Cost of not paying	Cost of paying, expected	Paying less not paying	Cheaper on this measure
20	16.00	13.76	(2.24)	Paying

Fictional. Computed in Python. Not paying: days x 0.8. Paying: 4.0 + 0.6 x 5 x 0.8 + 0.4 x (3 + days) x 0.8. At 9 days: 7.20 against 4.0 + 2.40 + 3.84 = 10.24. Break even at 15.33 days. Costs common to both paths, such as forensics and legal fees, are excluded. Not a recommendation to pay or not to pay.

What the framework shows. At nine days of restoration, which is the fictional company's current capability, paying is 3.04 million dollars more expensive on this measure, and even a decryptor that worked with certainty would cost 8.0 million against 7.2 million for restoring. Paying becomes cheaper only if restoration would take more than about 15.3 days, which is why the investment in tested, immutable backups of control O1 changes the position before any incident occurs. The ransom that would make paying equal to not paying at nine days is 0.96 million, a figure that shows how far the demand is from the economics. The analysis leaves out what may matter most: the legal position, the possible theft of data that encryption does not undo, the risk of repeated demands and the effect on customers, and each of these belongs in the advice of counsel and the forensic team. The CFO's contribution is to have the cost per day of an outage, the restore time and the insurance position known before the day on which the demand arrives (Module 12).

F. Insurance. Cyber insurance transfers part of the loss above a retention, the amount the company bears first, up to a limit. Its value depends on terms that the CFO should read with the broker and counsel, since policies commonly contain exclusions, conditions on controls such as multifactor authentication, sublimits on categories of cost and requirements to notify the insurer promptly, and these vary by insurer and year. Exhibit 18 applies a fictional policy with a retention of 1.0 million dollars, a limit of 15.0 million and a premium of 1.1 million to four loss sizes. The policy leaves a severe event of 32.0 million with a retained loss of 17.0 million, which is inside the tolerance of 22.5 million, while the uninsured loss would be above it.

Exhibit 18 Effect of a Cyber Policy with a 15.0 Limit and a 1.0 Retention (Millions of Dollars)

Loss event	Total loss	Recovered from insurer	Retained by company	Retained against tolerance of 22.5
Small incident	1.2	0.2	1.0	Within
Moderate incident	4.5	3.5	1.0	Within
Major ransomware event of R1	18.0	15.0	3.0	Within
Severe event	32.0	15.0	17.0	Within

Fictional. Computed in Python: recovered equals the loss less the retention, capped at the limit. 18.0 less 1.0 is 17.0, capped at 15.0, so 3.0 is retained; 32.0 less 15.0 leaves 17.0. At the register probability of 0.10, the expected recovery on R1 is 0.10 x 15.0 = 1.50 against a premium of 1.10 that covers all events. Exclusions, sublimits and conditions are not modeled.

G. Model drift, forecast error and fairness. The second line tests models with a small set of measures. The first measures drift in the data. The **population stability index** compares the distribution of an input or a score now with its distribution when the model was approved, by summing for each band the difference in share multiplied by the natural logarithm of the ratio of shares. A common rule of thumb, which the author states from general practice and not from a source opened for this module, treats a value below 0.10 as stable, from 0.10 to under 0.25 as a change to investigate and 0.25 or more as a significant shift. Exhibit 19 applies it to a fictional score distribution that is drifting over three quarters.

Exhibit 19 **Population Stability Index for a Fictional Scoring Model, Shares in Percent**

Score band (share of population)	Baseline	Quarter 1	Quarter 2	Quarter 3
Band 1 (lowest)	20	18	14	8
Band 2	20	19	17	14
Band 3	20	20	20	20
Band 4	20	21	23	26
Band 5 (highest)	20	22	26	32
Population stability index		**0.005**	**0.046**	**0.203**
Reading		Stable	Stable	Investigate

Fictional. Computed in Python. For quarter 3 the sum of (share less 0.20) x ln(share / 0.20) over five bands is 0.1100 + 0.0214 + 0.0000 + 0.0157 + 0.0564 = 0.2035. The thresholds are a common rule of thumb and not a regulatory standard.

Outcomes analysis. The second test compares forecasts with actual results, which SR 11-7 calls outcomes analysis. Exhibit 20 shows six months of a fictional revenue forecast produced with the help of a model. Across the six months the mean absolute percentage error is 6.78 percent, which is inside the limit of 8 percent set by the model owner, and the mean error is +6.0, which shows that the model is biased high in every month except one. The first three months averaged 3.58 percent and the last three 9.97 percent, and the rolling measure breaches the limit even though the full period passes. A test applied to an average of the whole period would have missed the deterioration.

Exhibit 20 **Backtest of a Fictional Forecast Model (Millions of Dollars)**

Month	Forecast	Actual	Error	Absolute percentage error
Month 1	102.0	100.0	+2.0	2.00 percent
Month 2	98.0	101.0	-3.0	2.97 percent
Month 3	110.0	104.0	+6.0	5.77 percent
Month 4	105.0	99.0	+6.0	6.06 percent
Month 5	120.0	108.0	+12.0	11.11 percent
Month 6	115.0	102.0	+13.0	12.75 percent
Months 1 to 3, average			+1.67	3.58 percent
Months 4 to 6, average			+10.33	9.97 percent (above the 8 percent limit)
All six months			**+6.00**	**6.78 percent**

Fictional. Computed in Python: absolute errors divided by actuals are 2.00, 2.97, 5.77, 6.06, 11.11 and 12.75 percent, whose average is 6.78 percent; the last three average 9.97 percent. The mean error is 6.00.

Fairness testing. The third test compares how a model treats groups. A simple measure is the selection rate, the share of applicants who are approved, and the ratio of the lower rate to the higher. In United States employment practice a ratio below four fifths has long been used as a signal that warrants review, and this is a screening rule of thumb and not a legal conclusion, which the author states from general knowledge. Which groups may be compared, what the law requires and what remedies are available are questions for counsel. Exhibit 21 compares a fictional customer credit process before and after an AI recommendation tool was introduced.

Exhibit 21 **Selection Rate Comparison for a Fictional Credit Process**

Process	Group A approved	Group A rate	Group B approved	Group B rate	Ratio of B to A	Signal
Manual review, prior year	450 of 800	56.2 percent	365 of 700	52.1 percent	0.927	No signal
With AI recommendation	480 of 800	60.0 percent	330 of 700	47.1 percent	0.786	Below 0.8, review

Fictional. Computed in Python: 480 / 800 = 0.6000, 330 / 700 = 0.4714, ratio 0.7857; 450 / 800 = 0.5625, 365 / 700 = 0.5214, ratio 0.9270. A ratio below 0.8 is a screening signal and not a legal finding. Consult counsel.

H. Scoring an inventory of AI use cases. An inventory lists every AI use in the company, and a scoring rule sorts the entries into tiers so that effort is spent where harm could be greatest. The method rates each use case from 1 to 5 on six dimensions, weights them, and maps the weighted score to a tier, with an override that places any use case with a top score on autonomy or on regulatory exposure in the highest tier. The weights in Exhibit 22 are 20 percent for data sensitivity, 25 for decision impact, 20 for autonomy, 15 for external exposure, 10 for regulatory exposure and 10 for the need for explainability, which sum to 100 percent. Tier 1 is a weighted score below 2.0, Tier 2 is 2.0 to below 3.0, and Tier 3 is 3.0 or more.

Exhibit 22 **Scored Inventory of Ten AI Use Cases**

No.	Use case (fictional company)	Data	Impact	Auton.	Exposure	Regul.	Explain	Weighted score	Tier
U1	Summaries of internal meetings	2	1	1	1	1	1	1.20	1
U2	Drafts of collection emails, sent after staff review	3	2	2	4	2	1	2.40	2
U3	Invoice coding suggestions, accepted by an accountant	3	3	2	1	2	2	2.30	2
U4	Draft commentary for the board package	4	3	2	2	2	2	2.65	2
U5	Anomaly detection in journal entries	4	3	2	1	3	3	2.70	2
U6	Customer credit limit recommendations	4	4	3	3	4	5	3.75	3
U7	Automated route price quotes to customers	3	4	4	5	2	3	3.65	3
U8	Agent that releases vendor payments and customer credits	5	5	5	2	3	3	4.15	3
U9	Screening of job applicants	4	4	3	3	5	4	3.75	3
U10	Monitoring of expense reports for policy breaches	3	3	1	1	5	3	2.50	3 (override)
	Weights	0.20	0.25	0.20	0.15	0.10	0.10	**1.00**	

Fictional. Computed in Python. Example U8: 5 x 0.20 + 5 x 0.25 + 5 x 0.20 + 2 x 0.15 + 3 x 0.10 + 3 x 0.10 = 4.15. U10 scores 2.50 and is raised to Tier 3 because its regulatory exposure is 5. Tier counts: Tier 1, one use case; Tier 2, four; Tier 3, five. Average weighted score 2.905.

Reading the inventory. Half of the use cases fall in the highest tier, which is a realistic result once the company counts agents, customer facing tools and tools that touch employees. The override matters: the monitoring of expense reports scores 2.50, below the Tier 3 line, because it is not autonomous and its data are moderately sensitive, but it touches employment and so brings legal exposure, and the rule forces a closer look. A weighted score is a screening device and not a measurement of risk, and the weights should be reviewed by the committee, but the discipline of scoring every use on the same dimensions makes the discussion concrete and exposes the use cases that nobody had thought to list. Exhibit 23 states what each tier requires.

Exhibit 23 Requirements by Tier in the AI Governance Operating Model

Requirement	Tier 1: standard (below 2.0)	Tier 2: enhanced (2.0 to below 3.0)	Tier 3: full review (3.0 or more, or override)
Approval	Owner records the use in the inventory	Owner and information security approve	AI governance committee approves, with legal, privacy and finance
Testing before use	Basic check of outputs	Documented testing on representative cases	Independent validation, bias test where people are affected, red team test for injection
Human oversight	Users review what they use	Named reviewer, sampling of accepted outputs	Human approval above stated limits, stop switch, no unreviewed action on people
Data	Approved tools only, no confidential data	Data classes approved, retention terms checked	Privacy review, contract terms on training and retention, data minimization
Monitoring	Annual attestation	Quarterly review of errors and overrides	Monthly indicators, drift and fairness tests, quarterly report to the committee
Incident handling	Report through the standard channel	Report within one business day	Immediate escalation, rehearsed playbook, vendor and regulator contacts

Source: author's proposal for teaching, informed by the Govern, Map, Measure and Manage functions of the NIST AI RMF. The tiers and cutoffs are examples and are not a regulatory standard.

I. The operating model. A policy that is written and not operated protects no one. The AI governance operating model has five parts that run in a cycle: a **policy** that states what is permitted and prohibited, an **inventory** that records every use, an **approval** process proportionate to tier, **monitoring** that keeps testing after launch, and **incident response** for the day something goes wrong. The committee that runs it should include finance, legal, privacy, security, technology and the business, and should report to the audit committee or risk committee of the board (Module 9). Exhibit 24 sets out the parts, the owner and the cadence.

Exhibit 24 The Operating Model for AI Governance: Policy, Inventory, Approval, Monitoring and Incident Response

Part	What it contains	Owner	Cadence and evidence
Policy	Permitted and prohibited uses, data rules, human oversight rules, vendor rules, consequences	General counsel and CFO	Annual review, board approval, signed acknowledgement by users

Part	What it contains	Owner	Cadence and evidence
Inventory	Every AI use, vendor, data class, owner, tier and approval date, including tools embedded in software and tools staff use without approval	AI governance lead	Continuous, with a quarterly sweep of network and expense data for unlisted tools
Approval	Scoring as in the inventory exhibit, tests required by tier, sign offs, conditions and limits	Committee chair	Before launch, and again on any change of model, data or purpose
Monitoring	Key risk indicators, error and override rates, drift and bias tests, cost against value, vendor changes	Use case owner, second line	Monthly or quarterly by tier, results to the committee
Incident response	Definition of an AI incident, escalation, containment by stop switch, evidence preservation, legal and disclosure review, lessons learned	CISO and general counsel	Rehearsed twice a year, with a record of each incident and its closure

Source: author's design for teaching. The parts follow the logic of the NIST AI RMF and the three lines model, and the cadence is an example.

J. Third party concentration and exit. Modern companies run on suppliers, and the CFO should know how much of the critical activity of the company depends on each. The **Herfindahl index**, the sum of the squared percentage shares, is a standard measure of concentration, and it rises sharply when one supplier dominates. In Exhibit 25 the primary cloud provider carries 46 percent of the critical processes of the fictional company, and the index is 2,912. Moving 15 points of the load to a second provider, as control O4 proposes, lowers the index to 1,982. The index has no regulatory threshold in this use, and it is helpful chiefly for comparing the position before and after a decision.

Exhibit 25 Concentration of Critical Processes by Provider (Fictional)

Provider	Share of critical processes, now	Share after second region	Squared share, now	Squared share, after
Primary cloud provider	46 percent	31 percent	2,116	961
Payments processor	22 percent	22 percent	484	484
Enterprise resource planning vendor	14 percent	14 percent	196	196
Carrier network platform	8 percent	8 percent	64	64
Identity provider	6 percent	6 percent	36	36
All others	4 percent	4 percent	16	16
Second cloud provider (new)	0 percent	15 percent	0	225
Concentration index (sum of squares)	**100 percent**	**100 percent**	**2,912**	**1,982**

Fictional. Computed in Python: 46 squared is 2,116, 22 squared is 484, 14 squared is 196, 8 squared is 64, 6 squared is 36, 4 squared is 16, sum 2,912. After: 31 squared is 961, plus 484, 196, 64, 36, 16 and 225, sum 1,982.

Exit plans. A contract that cannot be left is a dependence, and a company should know how long a replacement would take and how much help the contract promises while it moves. Exhibit 26 compares, for six suppliers, the months needed to move with the months of transition assistance that the contract provides. In 5 of the 6 cases the exit takes longer than the assistance, and for the cloud provider the gap is 8 months, which means that the company would have to run unsupported, or pay for an extension at a price set by the provider. Fourth parties deserve the same question: the CFO should ask each critical vendor which subcontractors hold the company's data or run its service, and whether the same few providers sit beneath

many of them.

Exhibit 26 Exit Time Against Contractual Transition Assistance (Fictional)

Supplier	Criticality	Months to exit	Months of transition assistance	Gap in months
Primary cloud provider	Core	14	6	8
Payments processor	Core	9	6	3
Enterprise resource planning vendor	Core	18	12	6
Carrier network platform	Important	4	3	1
Identity provider	Core	6	6	0
AI model provider, programming interface	Important	3	1	2

Fictional. Gap equals months to exit less months of assistance. Five of the six suppliers show a positive gap: 8, 3, 6, 1 and 2 months. The identity provider shows none.

Contract clauses. The terms of the contract are the main protection that the CFO controls directly, and they are best negotiated at the start when the supplier wants the business. Exhibit 27 lists the clauses that matter most for security, privacy and AI, with the reason for each. Counsel should draft and review them, and the right balance depends on the size of the supplier and the importance of the service.

Exhibit 27 Contract Clauses for Security, Privacy, Exit and AI Terms

Clause	What to require	Reason
Security standards and evidence	Named security controls, right to receive current independent assessment reports, and prompt remediation of findings	Allows the company to verify the state of the supplier and not rely on assertion
Incident notice	Notice of any incident affecting company data within a stated short period, with cooperation in investigation	Starts the clock for the company's own decisions, including disclosure analysis
Subcontractors and fourth parties	List of subprocessors, notice of change, flow down of the same obligations	Reveals concentration beneath the supplier
Data use, retention and return	No use of company data to train models without consent, defined retention, deletion and return at the end	Controls leakage and intellectual property exposure
Audit and testing rights	Right to audit, or to receive audit results, and to test continuity	Supports the second and third lines
Liability and insurance	Liability caps consistent with the risk, and supplier cyber insurance	Aligns the supplier with the cost of its failure
Exit and transition	Transition assistance for the time needed to move, data in a usable format, and no penalty for assistance	Closes the gaps measured in the exit table
AI specific terms	Disclosure of when AI is used in the service, notice of model changes, rights over outputs, indemnity for infringement	Addresses the new risks that older contracts do not mention

Source: author's synthesis from general practice. Contract terms should be drafted and reviewed by qualified counsel.

K. Business continuity. Continuity planning is the part of risk management that proves recovery, and it is treated as a discipline in Module 12. The CFO's interest is in the measures that connect an outage to cash: the

recovery time objective, the longest outage the business can bear for a process, and the demonstrated time in the latest test. In Exhibit 28 three of the five processes missed their objective in the latest test, including dispatch, which stops revenue, and treasury payments. A stated objective that has never been met is a hope and not a control, and the board should see the results of tests and not only the plan.

Exhibit 28 Recovery Objectives Against Demonstrated Recovery (Fictional)

Process	Recovery time objective, hours	Demonstrated in last test, hours	Result	Margin in hours
Dispatch and routing	8	14	Missed	-6
Customer billing	48	30	Met	+18
Payroll	72	60	Met	+12
Customer portal	12	20	Missed	-8
Treasury payments	24	26	Missed	-2

Fictional. Result is Met if the demonstrated time is no more than the objective. Dispatch 14 against 8, portal 20 against 12 and treasury payments 26 against 24 missed; billing 30 against 48 and payroll 60 against 72 met. Margin is the objective less the demonstrated time.

4. Executive case study and decisions

The case (fictional). The management of Harbourline Freight Systems, the fictional company of this module, asks the board to approve an artificial intelligence agent that reads supplier invoices and customer disputes and then releases vendor payments and issues customer credits on its own. The finance operations team handles 50,000 such items each year, valued at 109.5 million dollars in all, and a manual item costs 12 dollars to process against 1.50 dollars for the agent, a saving of 525 thousand dollars a year. The chief financial officer is asked to confirm that the proposal fits the risk appetite and to recommend how much freedom the agent should have. The inventory scoring of Exhibit 22 places this use case, U8, at 4.15, the highest in the company, and the committee of Tier 3 therefore reviews it. Exhibit 29 gives the facts. The numbers were computed in Python and are in thousands of dollars.

Exhibit 29 Profile of the Items the Agent Would Handle (Fictional, Thousands of Dollars)

Item	Count	Average value	Total value (thousands)	Error rate	Expected gross error (thousands)
Small, under 1,000	30,000	400	12,000	0.5 percent	60.00
Medium, 1,000 to 5,000	15,000	2,500	37,500	0.5 percent	187.50
Large, over 5,000	5,000	12,000	60,000	0.5 percent	300.00
Total	**50,000**		**109,500**	**0.5 percent**	**547.50**

Fictional. Computed in Python: 30,000 x 400 = 12,000; 15,000 x 2,500 = 37,500; 5,000 x 12,000 = 60,000; total 109,500. Gross expected error is 0.5 percent of value, 547.50. Assumptions: 70 percent of erroneous amounts are never recovered; a human reviewer catches 90 percent of errors in items reviewed at a cost of 15 dollars an item.

The decision point. The agent is accurate enough to be useful, with an error rate of 0.5 percent in the pilot, but two features worry the CFO. Errors are costly in the large items, which are 10 percent of the count and 55 percent of the value. And the agent reads text from outside the company, so a crafted invoice could carry

hidden instructions of the kind described as prompt injection. The security team estimates a 10 percent chance each year that such an attempt succeeds, with a loss of 2,190 thousand dollars, which is 2 percent of annual value, if the agent has no limits, and a loss of about 500 thousand dollars if a daily aggregate cap of 250 thousand dollars is in place and detection takes two days. Exhibit 30 compares five options.

Exhibit 30 Five Options for the Payment Agent (Fictional, Thousands of Dollars per Year)

Option	Items reviewed by a person	Saving	Expected error loss	Review cost	Expected injection loss	Net annual value
A. Fully autonomous, no limits	0	525.0	383.2	0.0	219.0	(77.2)
B. Autonomous with a daily aggregate cap	0	525.0	383.2	0.0	50.0	91.8
C. Cap and human approval of large items	5,000	525.0	194.2	75.0	50.0	205.8
D. Cap and approval of medium and large items	20,000	525.0	76.1	300.0	50.0	98.9
E. Do not deploy	Manual	0.0	0.0	0.0	0.0	0.0

Fictional. Computed in Python. Net value equals saving less expected error loss, review cost and expected injection loss. Option C: 525.0 less 194.25 less 75.0 less 50.0 is 205.75. Option D adds 225.0 of review cost to remove only 118.1 of expected loss. Option A: 525.0 less 383.25 less 219.0 is (77.25). Assumes 70 percent of errors are unrecovered.

The recommended path. The CFO recommends Option C: the agent may act alone on items of 5,000 dollars or less, a person approves every item above that amount, and a daily aggregate cap and a stop switch are in place from the first day. The expected net value is 205.8 thousand dollars a year, the best of the five, and well above the 91.8 thousand dollars of the cap alone. The marginal analysis explains the choice. Reviewing the 5,000 large items costs 75.0 thousand dollars and removes 189.0 thousand dollars of expected loss, a return of 2.5 to 1. Extending review to the 15,000 medium items costs a further 225.0 thousand dollars and removes only 118.1 thousand dollars, so the extra review costs more than it saves. Option A has a negative expected value, and the manipulation loss is the reason: without a cap the agent becomes the weakest point of the payment process. The recommendation carries conditions, including a pilot on a subset, monthly reporting of overrides and errors to the committee, vendor terms on model changes, and a review of the limits after six months.

What a weaker answer misses. A weaker answer would take the saving of 525 thousand dollars, note that the error rate is only half of one percent, and approve full autonomy, which on these figures loses money. It would treat human review as a uniform cost to be minimized and not as a control to be placed where it removes the most loss. It would overlook the hidden instructions in outside text, which the scoring of the use case should have warned of, and would make no provision for the case in which the model changes or the vendor retires it. It would also leave the record incomplete: no inventory entry, no tier, no owner and no stated indicators, so that a failure two years later could not be explained to the audit committee or, if the incident became material, to regulators. A stronger answer ties the proposal to the appetite, prices the alternatives, places the people and limits where they matter, and writes down the conditions under which the board would stop the program.

The wider lesson. The analysis does not depend on the particular figures. It shows that the economics of an autonomous system are set less by its average accuracy than by the size of its worst errors and by the controls around it, which is why the CFO, who understands both the value at stake and the ledger that records it, belongs at the center of AI governance. The same pattern, with a limit on the size of any action and a person at the points of greatest consequence, applies to systems that send customer messages, change prices or

approve credit, and the pattern is the practical meaning of the phrase human in the loop. Human oversight is useful only if the person has the authority, the information and the time to disagree, and a reviewer who approves every item in seconds provides the appearance of control and little more. The monitoring of the approval rate is therefore as important as the approval itself: a reviewer who accepts every item, or a rate of override that falls to nothing, is a signal of over reliance and should itself be an indicator on the dashboard.

5. Boardroom discussion and key takeaways

Directors do not need to understand how a model works to oversee it. They need to know that management has an inventory, that the risk is measured in dollars against a stated appetite, that the controls have owners and evidence, and that an incident would be recognized, escalated and disclosed in time. The questions in Exhibit 31 can be asked without technical knowledge, and the answers can be prepared from the exhibits of this module. Module 9 describes how the audit committee organizes this oversight.

Exhibit 31 Ten Questions Directors Should Ask About Enterprise, Cyber and AI Risk

Question a director should ask	What a ready CFO can say
What is our risk appetite in dollars, and are we inside it?	Appetite and tolerance statements, a register with expected losses and the test of Exhibit 14, with the excess and the plan to close it
Which three risks matter most, by expected loss and not by color?	The ranking by expected loss next to the ranking by score, with the reasons for any difference
Where does AI run in the company, and who approved it?	The inventory with tiers and the count of uses without approval, which should be zero (Exhibit 22)
What can an AI system do without a person, and what are its limits?	The permissions, amount and daily limits and stop switch for each Tier 3 use, and the override rate of reviewers
How would we learn of a cyber incident, and who decides that it is material?	The time to detect, the materiality procedure, the four business day clock and the Item 1.05 template (Exhibit 9)
If systems were encrypted tomorrow, what would we do about a ransom demand?	The playbook, the cost per day of outage, the tested restore time, counsel and insurer roles, and the break even analysis (Exhibit 17)
Which suppliers could stop us, and how long would it take to leave?	The concentration measure, the exit gaps and the plan to close them (Exhibit 26)
Have our recovery times been proven in a test?	Objectives against demonstrated results, and the schedule of corrective action (Exhibit 28, Module 12)
How do we know our models still work and treat people fairly?	Drift, error and fairness measures reviewed by the second line, with the thresholds that trigger a review
What evidence would an auditor or a regulator want, and do we have it?	Records of approvals, tests, incidents and board reports, and the link to internal control and assurance (Module 11)

Source: author's synthesis. Cross references are to modules in this program: Module 9 for governance, Module 11 for reporting and assurance and Module 12 for crisis and resilience.

Advising the CEO and the board. The CFO should make three arguments. The first is that risk should be stated in dollars and compared with a limit that the board has chosen, because a color on a map gives no basis for a decision on how much to spend. The second is that the controls for new technology should be placed at the points of greatest consequence, and that autonomy should be granted in steps as evidence accumulates. The third is that the evidence for these controls is a record kept every day, including the record that explains why a risk was accepted. PwC reported in October 2025, from 3,887 business and technology leaders in 72

countries, that 60 percent placed cybersecurity investment among their top three strategic priorities, while only 6 percent felt fully prepared for all the vulnerabilities listed, which suggests that spending and readiness are not the same thing.

Common errors. Eight errors recur. The first is to rank risks by multiplied scores and to ignore expected loss and tolerance. The second is to write the appetite in words that cannot be tested. The third is to treat cyber as the concern of the technology department alone. The fourth is to keep no inventory of AI and to discover shadow use after an incident. The fifth is to grant autonomy to a system on the strength of its average accuracy and not of its worst error. The sixth is to buy insurance without reading its conditions, or to rely on it as the whole plan. The seventh is to sign contracts with suppliers that cannot be left. The eighth is to plan for recovery and never to test it.

Key takeaways. First, enterprise risk management is useful only when it compares risk with an appetite that the board has set in numbers. Second, compute expected loss alongside the score, and state the tolerance for a single event. Third, scores hide the difference between frequent small losses and rare large ones, so the CFO should read the register in dollars. Fourth, model risk is controlled through validation, monitoring and effective challenge by the second and third lines. Fifth, the risks of AI, hallucination, injection, leakage, intellectual property, bias, over reliance and agentic action, are answered by limits, approvals and data rules and not by trust in the model. Sixth, use the NIST frameworks as checklists, and confirm the current text of the EU AI Act and the SEC rules with counsel. Seventh, know the cost of an incident bottom up, understand the limits of insurance, and decide on the framework for a ransom demand before the day it arrives. Eighth, measure concentration among suppliers, close the exit gaps and negotiate the clauses that protect the company. Ninth, prove recovery by test. Tenth, run AI governance as a cycle of policy, inventory, approval, monitoring and incident response, with the audit committee in sight.

6. The executive artifact: a risk register extract and an AI use case approval checklist

The first artifact is an extract of the enterprise risk register, in the form that a committee reviews each quarter. It carries what the committee needs in order to act: the risk event, the owner, the current rating and expected loss, the controls, the indicator that gives warning, the treatment decision and the date of the next review. Exhibit 32 shows four entries completed for the fictional company, using figures computed in Section 3. The reader should prepare a blank register with the company's own risks, have the legal and insurance entries reviewed by counsel and the broker, and file each version so that the history of every decision can be reconstructed.

Exhibit 32 Enterprise Risk Register Extract, Completed for the Fictional Harbourline Freight Systems

Field	R1 Ransomware	R6 Public AI tool leakage	R3 Cloud outage	R7 Largest customer
Owner	CIO	CISO	CIO	CEO
Probability, impact, expected loss	10%, 18.0, 1.80	45%, 2.5, 1.12	8%, 12.0, 0.96	7%, 35.0, 2.45
Score and rating	8, Moderate	8, Moderate	8, Moderate	10, High
Controls in place or planned	O1 immutable backups and segmentation, cyber policy with 15.0 limit	O3 approved AI tool with data loss prevention, user policy	O4 second region and tested failover	Contract renewal program, pricing and service review

Field	R1 Ransomware	R6 Public AI tool leakage	R3 Cloud outage	R7 Largest customer
Expected loss after planned controls	0.60	0.38	0.32	2.45 (unchanged)
Key risk indicator and trigger	Restore tests meeting target, Red below 90 percent	Traffic to unsanctioned tools, Red above 10 percent	Demonstrated recovery against objective, Red if missed	Share of revenue under contract beyond 12 months
Treatment decision	Reduce and transfer	Reduce	Reduce, required to meet appetite	Accept for now, escalate to the board as above tolerance
Next review	Quarterly, committee	Monthly until indicator is Green	After failover test	Quarterly, board

Fictional. Figures from the register and control exhibits of Section 3 (millions of dollars): R1 0.10 x 18.0 = 1.80 and 0.05 x 12.0 = 0.60 after O1; R6 0.45 x 2.5 = 1.125 and 0.15 x 2.5 = 0.375 after O3; R3 0.08 x 12.0 = 0.96 and 0.08 x 4.0 = 0.32 after O4; R7 0.07 x 35.0 = 2.45. The extract is an illustration and a proposed program standard.

The AI use case approval checklist. The second artifact is the checklist that the owner of a use case completes, and that the committee reviews, before an AI system is approved for use. It follows the order of the NIST functions, from the context of the use to the evidence of monitoring, and its fields correspond to the tests of Section 3. Exhibit 33 shows the checklist completed for the payment agent of the case, use case U8, as it stood at the time of the recommendation.

Exhibit 33 AI Use Case Approval Checklist, Completed for the Payment and Credit Agent (U8)

Item	Requirement	Entry for U8, the payment and credit agent (fictional)	Status
1. Purpose and owner	Written purpose, named business owner and technical owner	Release of vendor payments and customer credits. Owner: controller. Technical owner: head of applications	Complete
2. Inventory and tier	Entered in the inventory, scored, tier assigned	Weighted score 4.15, Tier 3 (Exhibit 22)	Complete
3. Data	Data classes listed, privacy review, vendor terms on retention and training	Supplier and customer data. Privacy review open. Vendor terms prohibit training on company data	Partial
4. Testing	Accuracy on representative cases, bias test where people are affected, test for prompt injection	Pilot error rate 0.5 percent. No people affected. Injection test by security team scheduled before pilot	Partial
5. Human oversight	Approval above stated limits, reviewer authority and capacity	Approval of items above 5,000 dollars, 5,000 items a year, two trained reviewers	Complete
6. Limits and stop switch	Per item and daily caps, ability to stop within minutes	Daily aggregate cap of 250 thousand dollars, stop switch held by the controller and CISO	Complete
7. Monitoring	Indicators, thresholds, reviewer override rate, drift and cost tracking	Monthly report of errors, overrides and approval rate to the committee, population stability review	Complete

Item	Requirement	Entry for U8, the payment and credit agent (fictional)	Status
8. Vendor and exit	Contract clauses, notice of model change, exit plan	Notice of model change required. Exit plan for the model provider, 3 months, assistance 1 month	Gap
9. Incident response	Defined incident, escalation, rehearsal, disclosure review	Included in the AI incident playbook. First rehearsal planned before launch	Partial
10. Legal and regulatory	Counsel review of applicable law and contracts, EU AI Act relevance	Counsel review in progress. No use of outputs in filings	Partial
11. Approval and review date	Committee decision, conditions, date of next review	Approve Option C subject to items 3, 4, 8, 9 and 10. Review after six months	Conditional

Fictional. The entries reflect Option C of Section 4. Status: Complete 5, Partial 4, Gap 1, Conditional 1. Counsel should review the legal and regulatory entries, and the checklist is a proposed program standard.

Using the artifacts. The register extract should be read together with the key risk indicators and the tolerance, and an entry in which the expected loss is above the appetite, or the single event above the tolerance, should be on the agenda of the board and not only of the committee. The checklist should be completed by the owner and not by the governance function, so that ownership is real, and an item marked Partial or Gap should carry a date and a person. In the example, five items are complete and the approval is conditional on the rest, which is a better position than approval on the strength of enthusiasm, and a record that will be useful if an incident or an audit follows.

7. Research and Perspectives from the Leading Firms

The publications below were found by search, and each page was opened before it was described. Where a page returned only part of its content, or did not state its date, the description says so and states only what was returned. The expected loss arithmetic, the scoring of risks and use cases, the decision frameworks and the drift and fairness measures in Sections 2 and 3 are the author's own and not taken from these pages. Rules and practice change quickly in this field, and the reader should confirm the current text before relying on any item.

Exhibit 34 Publications and Standards Used in This Module

Source	Title and URL	What it says
IBM	Cost of a Data Breach Report 2026 (page opened October 2026). https://www.ibm.com/reports/data-breach	Reports a global average cost of a breach of 4.99 million dollars, up 12 percent and a record, a saving of 1.93 million dollars for organizations that use security AI and automation extensively, and a 56 percent increase in AI driven attacks. The page did not state the United States average or the time to contain
PwC	2026 Global Digital Trust Insights (October 20, 2025), Dutch edition page. https://www.pwc.nl/en/insights-and-publications/themes/digitalization/2026-global-digital-trust-insights-survey.html	Survey of 3,887 business and technology leaders in 72 countries. 60 percent place cybersecurity investment in their top three priorities, 6 percent feel fully prepared for all listed vulnerabilities, and 49 percent have not acted on quantum risk. Third party breaches, cloud and connected products are leading themes

Source	Title and URL	What it says
EY	EY survey: AI adoption outpaces governance as risk awareness among the C-suite remains low, press release (August 13, 2025). https://www.ey.com/en_ro/newsroom/2025/08/ey-survey-ai-adoption-outpaces-governance-as-risk-awareness	Survey of 975 C-suite leaders of companies above 1 billion dollars of revenue in 21 countries, conducted in March and April 2025. 76 percent use or plan to use agentic AI within a year while 56 percent know its risks, and one third have the protocols for all facets of EY's responsible AI framework
Deloitte	The State of AI in the Enterprise, 2026 AI report (page undated, copyright 2026). https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html	Survey of 3,235 leaders in August and September 2025. Only one in five companies has a mature model for governing autonomous AI agents, while agentic use is expected to rise sharply in two years. States that senior leadership involvement in AI governance is linked to greater value, without a figure
KPMG	AI Quarterly Pulse Survey, Q3 2026 (report dated September 2026). https://kpmg.com/kpmg-us/content/dam/kpmg/pdf/2026/aipulsesurvey-q3.pdf	Survey of 314 United States leaders at organizations with revenue of 1 billion dollars or more, fielded July 24 to August 25. Reports that identification of high risk use cases where autonomous decisions are not allowed rose from 28 to 49 percent over a year, and lists the most common controls, such as data access controls at 59 percent
NIST	AI Risk Management Framework page and NIST AI 100-1. https://www.nist.gov/itl/ai-risk-management-framework and https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf	Voluntary framework released January 26, 2023, with the functions Govern, Map, Measure and Manage and seven characteristics of trustworthy AI. The page notes a revision under the White House AI Action Plan and the generative AI profile AI 600-1 of July 26, 2024
NIST	Cybersecurity Framework 2.0 (NIST CSWP 29, February 26, 2024) and release announcement. https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf	Defines six functions, adds Govern, and extends the framework beyond critical infrastructure to organizations of any sector
SEC	SEC Adopts Rules on Cybersecurity Risk Management, Strategy, Governance and Incident Disclosure, Press Release 2023-139 (July 26, 2023). https://www.sec.gov/newsroom/press-releases/2023-139	Requires Form 8-K Item 1.05 disclosure generally four business days after a determination of materiality, with a delay only on a written finding of the Attorney General, and annual disclosure under Item 106 of Regulation S-K
European Commission	AI Act, regulatory framework for AI (page opened October 2026). https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai	States entry into force on August 1, 2024, four risk tiers, application dates through August 2028 and the digital omnibus that entered into force on July 27, 2026 and delayed high risk rules. Gives no penalty amounts
COSO	Enterprise Risk Management guidance page. https://www.coso.org/guidance-erm	Lists the 2017 framework Integrating with Strategy and Performance, a 2021 paper on AI, a 2019 paper on cyber risk, a 2020 paper on risk appetite and a 2026 paper on practical ERM, without summarizing their content
ISO	ISO 31000:2018 Risk management, Guidelines. https://www.iso.org/standard/65694.html	Describes principles and guidelines for identifying, analyzing, evaluating, treating, monitoring and communicating risk, usable by any organization, not certifiable, confirmed in 2023 and marked to be revised

Source	Title and URL	What it says
Federal Reserve	SR 11-7, Guidance on Model Risk Management (April 4, 2011). https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm	Defines model risk, effective challenge and the three parts of validation, and scales the rigor to the size, nature and complexity of the organization

Each page was opened and the descriptions follow what the tool returned. No publication of RSM US, BDO, Grant Thornton, Baker Tilly, Crowe, Forvis Mazars, CBIZ, Moss Adams, CohnReznick, CLA or the PCAOB was read for this module, and none is cited. Survey statistics are self reported. The IBM page is a summary, and the full report was not read.

What the research agrees on. The sources differ in kind, from standards and rules to surveys of executives, and they should not be read as one body of evidence. They agree that the cost of cyber incidents is rising and that the third party is a principal route of attack, that the adoption of AI, and especially of agents that act, is running ahead of governance, and that the frameworks of NIST, COSO and ISO converge on a similar structure of accountability, context, measurement, action and recovery. None of them chooses an appetite or a budget, and these two decisions remain with the board and with finance.

How a CFO can use the research. The surveys are best used to calibrate expectations, such as how common a mature agent governance model is among peers, and the standards are best used as checklists. The SEC and EU pages are the most consequential, since they define obligations, and the CFO should confirm the present text with counsel. The IBM figure of 4.99 million dollars is a useful reference but a poor budget, and a bottom up estimate such as Exhibit 16 is a better basis for decisions.

8. The ten question self assessment

Rate each statement from 1 to 5 using the scale below, and mark one box in each row. The ratings must rest on actual experience or documented evidence, such as a risk register you maintained, a model you validated or an incident you managed, and not on confidence. A rating of 1 means Unprepared: limited understanding and no meaningful application. A rating of 2 means Developing: conceptual familiarity, relying heavily on others. A rating of 3 means Competent: able to apply the capability independently in conventional situations. A rating of 4 means Advanced: the capability is integrated into strategy and delivers measurable results. A rating of 5 means Future-ready: enterprise wide impact, anticipation of change and the development of others.

No.	Statement of demonstrated capability	1	2	3	4	5
1	I maintain a risk register with owners, scores and expected losses, and I compare it with a risk appetite and a tolerance that the board has approved in numbers.	☐	☐	☐	☐	☐
2	I can compute expected loss, read a heat map critically, and test a set of controls for cost and benefit against the appetite.	☐	☐	☐	☐	☐
3	I have defined key risk indicators with thresholds and owners, and I review them with the committee on a fixed cycle.	☐	☐	☐	☐	☐
4	I can explain model risk, effective challenge and the three lines model, and I have seen a model validated, monitored and tested for drift or bias.	☐	☐	☐	☐	☐
5	I can describe the risks of hallucination, prompt injection, data leakage, intellectual property, bias, over reliance and agentic action, and the controls that answer each.	☐	☐	☐	☐	☐

No.	Statement of demonstrated capability	1	2	3	4	5
6	I use the NIST AI Risk Management Framework and the NIST Cybersecurity Framework 2.0 as working checklists, and I know what the SEC rules and the EU AI Act require at a high level.	☐	☐	☐	☐	☐
7	I can estimate the cost of a cyber incident bottom up, I understand the terms and limits of insurance, and I have a decision framework for a ransom demand agreed in advance.	☐	☐	☐	☐	☐
8	I can describe the materiality process and the four business day disclosure clock for a cyber incident, and I know who decides and who must be informed.	☐	☐	☐	☐	☐
9	I measure supplier concentration, maintain exit plans, and negotiate security, privacy and AI clauses, including fourth party transparency.	☐	☐	☐	☐	☐
10	I run or sponsor an AI governance operating model with policy, inventory, tiered approval, monitoring and incident response, and I test continuity against stated recovery objectives.	☐	☐	☐	☐	☐

Total score: _____ out of 50. Ten statements rated from 1 to 5 give a minimum of 10 and a maximum of 50 points. Exhibit 35 gives the proposed interpretation. The scale is that of the source proposal for this program, and the statements are specific to this module.

Critical module note

Module 10 is one of the five critical modules in the program, together with AI literacy (Module 1), capital allocation (Module 5), governance (Module 9) and enterprise leadership (Module 13). A score below 30 out of 50 on a critical module would block the highest designation, even when the total across all twenty modules is high. The reason is that excellence in one area should not conceal a consequential weakness in another, and a weakness in risk is the kind that tends to be exposed at the worst moment.

Exhibit 35 Proposed Interpretation of the Module 10 Score

Score	Reading	Suggested use
10 to 19	Unprepared	Study the module in full, and rebuild the register, control and inventory examples with your own figures before the next risk review
20 to 29	Developing	Prepare a scored register with expected losses for one function, and an inventory of AI uses, and ask a peer to challenge both. A score in this range blocks the highest designation
30 to 39	Competent	Write the appetite and tolerance statements for board approval, and add key risk indicators and a materiality procedure for incidents
40 to 45	Advanced	Run the AI governance operating model, test recovery, and report to the audit committee with the exhibits of this module
46 to 50	Future-ready	Share the method across the enterprise, invite independent review of the register and of the governance record, and recalibrate probabilities from incidents and near misses

These bands are proposed program standards and have not been empirically validated against a broad population of finance executives. Module 10 is designated a critical module, and a score below 30 would block the highest designation.

Development actions by pattern. A single total can hide the pattern behind it, so look at which statements received the lowest ratings. Exhibit 36 suggests three actions for each pattern.

Exhibit 36 Development Actions for Low Score Patterns

Low scores on	Pattern	Three suggested development actions
Statements 1, 2 and 3	Measurement gap: the reader cannot yet state risk in dollars against an appetite	1. Rebuild Exhibit 11 with your own risks and compute expected loss. 2. Draft appetite and tolerance statements as in Exhibit 2. 3. Define five key risk indicators with thresholds as in Exhibit 15
Statements 4 and 5	Model and AI risk gap: the reader cannot yet control a model or an AI system	1. Compute the stability index and a backtest for a model you use as in Exhibit 19. 2. For each risk in Exhibit 5, name the control in your company. 3. Review a Tier 3 use case with the second line
Statements 6, 7 and 8	Cyber and disclosure gap: the reader cannot yet judge an incident or a disclosure duty	1. Walk the six functions of Exhibit 7 with the CISO. 2. Build a bottom up cost of an incident as in Exhibit 16 and read the insurance policy with the broker. 3. Draft a materiality procedure with counsel and rehearse it
Statements 9 and 10	Third party and operating model gap: the reader cannot yet govern suppliers or AI at scale	1. Compute concentration and exit gaps as in Exhibit 25 and Exhibit 26. 2. Stand up the inventory and the tiers of Exhibit 23. 3. Run a recovery test and report the demonstrated times as in Exhibit 28

Source: author's suggestions. The mapping of statements to patterns is the author's proposal, and ratings should rest on experience or documented evidence.

A closing thought. The purpose of risk management is to allow the company to take the risks that it wants with its eyes open, and the CFO who can state the appetite in dollars and has proved recovery by test can say yes to new technology with more confidence.

9. Currency check and sources

Currency check: this field changes, and statistics, rules, framework versions, insurance terms and survey results shown here may already have changed. The publications in Exhibit 34 were opened in October 2026, and their statistics are largely self reported. The following were confirmed on pages read: the IBM 2026 figures of 4.99 million dollars, 12 percent and 1.93 million dollars; the PwC survey statistics of October 20, 2025; the EY statistics of August 13, 2025; the Deloitte statistics on governance of agents, from the Deloitte page, whose publication date was not stated; the KPMG Q3 2026 pulse statistics; the NIST AI RMF dates and function descriptions; the NIST CSF 2.0 date and function descriptions; the SEC press release statements on Item 1.05, Item 106 and the effective dates; the European Commission dates for the AI Act and the omnibus; the COSO and ISO page statements; and the SR 11-7 definitions. A secondary news summary of the Deloitte survey that was also read gave the date as late January 2026, which is not relied on. The following come from the author's own knowledge and were not confirmed in a source read: the privacy law questions; the rule of thumb thresholds for the stability index and the four fifths ratio; the three lines model; the typical features of cyber insurance; the legal considerations in a ransom payment; the typical contract clauses; the Herfindahl index; and the scoring bands, weights, tiers and thresholds. The company and all figures in the exhibits are fictional, and all sums, expected losses, scores and tiers were computed in Python and checked by assertion. This module is educational and is not accounting, tax, legal, securities, privacy, insurance or security advice. Readers must consult qualified professionals in accounting, tax, legal and other relevant services before acting, and matters such as incident disclosure, ransom payments, privacy law, the EU AI Act, insurance coverage and contract terms in particular require the advice of qualified counsel, the insurance broker and security specialists. Every rate and assumption in the exhibits is an illustration, and a reader who applies the method must replace each with the company's own figure.

Sources: IBM Cost of a Data Breach Report 2026; PwC 2026 Global Digital Trust Insights (October 2025); EY Responsible AI Pulse press release (August 2025); Deloitte State of AI in the Enterprise; KPMG AI Quarterly Pulse Survey Q3 2026; NIST AI Risk Management Framework and NIST AI 100-1; NIST Cybersecurity Framework 2.0 (CSWP 29); SEC Press Release 2023-139; European Commission AI Act page; COSO ERM guidance page; ISO 31000:2018; Federal Reserve SR 11-7. Titles, dates and URLs appear in Exhibit 34. Harbourline Freight Systems and every figure in its exhibits are fictional.