WORKING DRAFT V0.11 · OPEN FOR REVIEW · NOVEMBER 2026

Business Outcome Risk Analysis for Enterprise AI (BORA-EAI)

Threats, value, and ownership for each AI use case. Driven by strategy and performance.
Enterprise AI risk is a puzzle of several pictures; BORA-EAI lines them up for the executives who own each one.

FIVE PICTURES, FIVE OWNERS, FIVE CURRENCIES

Law and regulation
General Counsel, HR
SETTLEMENTS AND FINES
EU AI Act: up to €15M or 3% of turnover for most deployer duties, €35M or 7% for prohibited practices.

Finance
CFO
BUDGET VARIANCE
Gartner expects at least 50% of GenAI projects to overrun budgets through 2028.

Operations
COO, revenue owners
REVENUE LOST, EFFICIENCY GAINED
Business interruption had the highest median loss, about US$90k per claim (Verizon 2026).

Security
CISO
INCIDENT AND BREACH COST
13% of organizations had an AI model or app breach; 97% of them lacked AI access controls (IBM 2025).

Strategy
CEO, board
GROWTH GAINED OR LOST
BCG's "future-built" firms show 2.4x revenue and 2.8x EBITDA growth vs laggards (a correlation).

WHERE EXISTING MODELS STOP

We reviewed more than a dozen frameworks, standards and guidance documents. Each covers a piece. None we found balances threats and benefits across all five owners.

- **Quantify loss:** FAIR-AIR
- **Tier by capability or map to controls:** CSA CBRA, NIST Cyber AI Profile (draft), AWS GenAI and agentic scoping matrices
- **Enumerate negative impacts:** Huwyler, MIT AI Risk Repository, Geneva Association, NIST AI 600-1
- **Say what to govern:** ISO 42001, NIST AI RMF, COSO, ISO 23894

Gaps we found: embedded AI in COTS products; social engineering of a chatbot's commercial logic; opportunity cost (lost revenue and back-office efficiency); internal governance cost; data sovereignty across deployment models. Even the closest fit, ISO 42001 with the NIST AI RMF, does not say which AI business scenarios to examine, who owns them, or how to price them.

HOW BORA-EAI WORKS

The use case is the unit. Who uses it, what data it touches, what it may do, and which executive carries the consequence. A coding assistant, a customer chatbot and a vendor-enabled AI feature are three use cases, even on one model, and each gets its own answer.

Both sides of the ledger. Benefits are weighed against losses and the cost of governing, as ISO 31000 and the NIST AI RMF (MAP 3.1 and 3.2) both allow. Declining to deploy is a cost, not a safe default.

Two lenses. *Internal:* permissions, data exposure, and whether the productivity gain is real. *Customer-facing:* what it may say or offer, who is liable, and what revenue it protects or puts at risk.

Depth matches stakes. A tier sets how deep the analysis goes (see the tiers below). Two override rules lift a use case to at least Medium: read access to regulated personal data, and customer-facing representations or concessions. Thresholds are our judgment, open to challenge.

DEPTH OF ANALYSIS BY TIER

MORE DEPTH →

Fast lane
Matches a pre-approved pattern (sanctioned tool, approved data class, no external actions or customer statements): registered, not assessed.

Low
Qualitative screen. No numbers.

Medium
Ranges for the top three scenarios and for value.

High
Full quantification, red teaming, insurance review, board reporting.

NINE LAYERS, ONE EXECUTIVE QUESTION EACH

	LAYER	EXECUTIVE QUESTION	DRAWS ON
1	**Use-case inventory**	What is AI doing for us, and where?	COSO capability types × lines of business
2	**Consumption and ownership**	Who controls the app, model and data? Who is accountable?	AWS GenAI scopes 1–5 (plus new 2e, vendor-enabled embedded AI); CSA AICM roles
3	**Sovereignty profile**	Where does it run, whose law applies, can we leave?	Provider residency tiers; APRA fourth-party framing
4	**Agency and tier**	How much can it do without us? Does it qualify for the fast lane?	AWS agentic scopes; CSA CBRA
5	**Business threat scenarios**	What could go wrong, in business terms?	New library, anchored to Huwyler and OWASP
6	**Loss**	What would it cost, and who pays? How bad is a bad year?	Geneva Association loss types; FAIR-AIR ranges, Monte Carlo
7	**Value**	What do we gain, or forgo by waiting?	ISO 31000; NIST AI RMF MAP 3.1–3.2; uplift studies
8	**Governance cost**	What does controlling it cost?	ISO 42001; EU AI Act cost estimates
9	**ERM fit and completeness**	Does it fit risk appetite? What did we miss?	ISO 23894; ISO 42001; ISO 42005; MIT repository

FIFTEEN BUSINESS THREAT SCENARIOS, ONE ACCOUNTABLE OWNER EACH

One executive is accountable for each use case's net position. Each scenario has one accountable owner (A), who holds the main lever that prevents the loss; others are consulted (C).

B1	**Insider amplification** **A** CISO · **C** CHRO	B2	**Sanctioned-AI leakage** **A** CIO (platform owner) · **C** CISO	B3	**Shadow-AI leakage** **A** CISO · **C** CIO, Procurement
B4	**Untrusted AI artifacts** Models, tools, MCP servers · **A** CTO · **C** CISO	B5	**IP and license exposure** Output copyright, licenses, indemnity · **A** General Counsel · **C** CTO, Procurement	B6	**AI vendor and fourth-party concentration** **A** CRO · **C** CIO, Procurement
B7	**Vendor-enabled embedded AI** **A** General Counsel · **C** DPO, Procurement	B8	**Bot misrepresentation liability** **A** Channel owner (CX) · **C** General Counsel	B9	**Bot technical compromise** **A** CISO · **C** Channel owner
B10	**Bot commercial manipulation** **A** Channel owner (CX, Sales) · **C** CFO	B11	**Agentic execution failure** Destructive actions, runaway spend, cascading errors · **A** CTO · **C** Continuity owner	B12	**Agentic unauthorized commitments** Payments, orders, contracts · **A** CFO · **C** General Counsel
B13	**Sovereignty, jurisdiction, continuity** **A** CRO · **C** General Counsel, CIO	B14	**Forgone and negative value** **A** Use-case sponsor (line head) · **C** CFO	B15	**Governance cost and failure** **A** CRO · **C** CFO

THE OUTPUT: NET POSITION PER USE CASE, AS A RANGE

Net annual position = P(reaches production) × P(delivers value | production) × gross annual value
– expected annual loss across B1 to B13 – allocated governance cost

Ranges, not points
Every input is a three-point estimate (low, likely, high), run through Monte Carlo, as FAIR does. Thin data means a wide range.

What the board sees
The 10th, 50th and 90th percentile net position, the chance it is below zero, and a 1-in-20-year loss per scenario.

Reading it
Uplift can fall below zero for experts or tasks outside a model's capability. B14 and B15 sit in the value and governance terms. The tier sets depth, never a loss figure.

STATUS AND HOW TO HELP

Not yet validated. New here: the scenario library, the three-way chatbot split, the sovereignty profile, the override rules and fast lane, and value-versus-governance netting; the rest is borrowed. B12 is forward-looking. **Tell us** what is missing, misowned or misworded, and what we have missed.