OpenQ

When Low Probability Becomes Unacceptable

Internal Auditor ในอุตสาหกรรมที่เปลี่ยนเร็วกว่าการประเมินความเสี่ยง

Clinic IA 5/2026 · IIAT × ISACA Bangkok Chapter · 27 June 2026

Noratip Dhanasarnsilp | CISA · CISSP · ITIL Expert

Session overview

01 **Mindset shift in 2025**
What happened and why it changes everything

02 **Hot topics and risks 2026**
Big4 research + Thai SET50 evidence

03 **The operational blind spots**
Where incidents really come from

04 **How IA responds**
Agile IA methodology + leveraging agentic AI

05 **Measuring maturity**
COBIT-style maturity model + 12-month roadmap

OpenQ

Mindset shift in 2025

01

What happened, and why it changes everything for IA

March 28, 2025.

A 7.7 magnitude earthquake.

- The Stock Exchange of Thailand suspended trading.
- The Bank of Thailand issued BCM guidance within 48 hours.
- Buildings across Bangkok evacuated.

Thailand was not classified as a seismic zone.

Were these surprises, or just uncomfortable to plan for?

OpenQ

July 19, 2024.

CrowdStrike released a routine content update to its Falcon security sensor at 04:09 UTC. Within 78 minutes, 8.5 million Windows machines were showing the Blue Screen of Death at the same time, across airlines, hospitals, banks, payment systems, and emergency services worldwide.

It was not a cyberattack.

Vendor concentration risk

One vendor · one update · 8.5M systems offline simultaneously. Your security solution was the disruption.

BCP gap

Recovery required manual intervention on every affected machine. No remote fix. BCPs did not cover this scenario.

Source: Microsoft post-incident report · CISA Alert July 19, 2024 · CrowdStrike PIR August 2024

One typo.

A routine database configuration change on a Tuesday afternoon. No second reviewer, no change window, treated as business as usual. A single wrong value takes a core system offline for hours.

The CrowdStrike outage began the same way: a routine content update.
The most damaging failures come from ordinary tasks, not exotic threats.

Familiarity lowers attention. Low attention on a high-impact task is exactly where catastrophic human error lives.

CrowdStrike was not a one-off. It is a category of risk

As organizations move to SaaS and shared infrastructure, one vendor failure becomes everyone's failure, at the same time, without warning.

July 2024 · Endpoint security vendor

CrowdStrike

Routine update crashes 8.5M Windows systems. Hospitals, airlines, banks, emergency services offline. $5.4B estimated impact on US Fortune 500.

⚠ *BCP gap: 'Security vendor IS the disruption.' Not in any BCP scenario library*

Nov 18, 2025 · CDN · WAF · DDoS protection

Cloudflare

Global outage takes down X, ChatGPT, Uber, and hundreds of business apps worldwide. Organizations lost security layer and availability simultaneously.

⚠ *BCP gap: WAF offline = unprotected AND unreachable. Two failures for the price of one*

Dec 5, 2025 · 18 days after the November outage

Cloudflare (again)

Second major outage. X, LinkedIn, Substack, Canva, Spotify, and Downdetector all down. Even the outage monitoring service was affected.

⚠ *BCP gap: Two critical outages from the same single-vendor dependency in one month*

June 2021 · Content delivery network

Fastly CDN

One-hour outage: Amazon, Reddit, CNN, NYT, UK Government websites offline simultaneously. ~$67M in lost commerce estimated.

⚠ *BCP gap: Concentration risk: 2% of the internet routed through one CDN provider*

The IA question: is 'a third-party vendor takes our critical infrastructure offline, at the same time, without warning, and it is not a cyberattack' in your risk register?

OpenQ

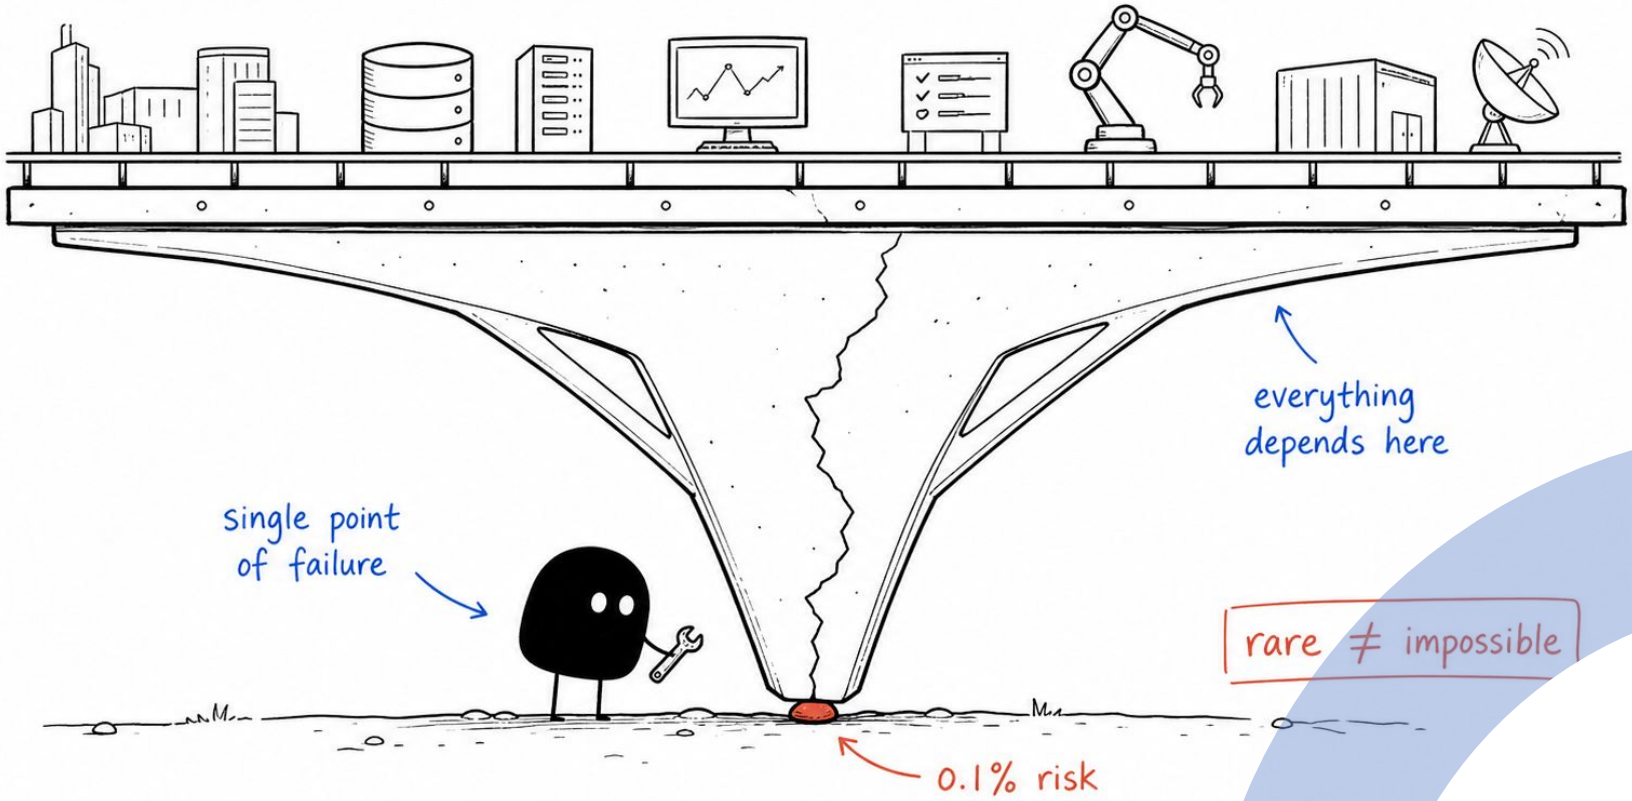
everything
depends here
single point
of failure
rare ≠ impossible
0.1% risk

Four events. One pattern.

→ Bangkok earthquake, March 2025. Thailand ‘not a seismic zone.’ SET suspended trading. BOT issued emergency BCM guidance.

→ CrowdStrike, July 2024. A routine security update crashed 8.5M systems at once. Hospitals, banks, airlines. Not a cyberattack.

→ Global conflict escalation. Energy price spikes across Asia. Labeled ‘low probability geopolitical tail risk’ in 2023 planning cycles.

→ COVID-19. A pandemic scenario sat in every BCP document. Almost no organization had run the drill when it was needed.

Common thread: none of these were rated ‘high probability’ on any risk heat map.

“Were these surprises, or just uncomfortable to plan for?”

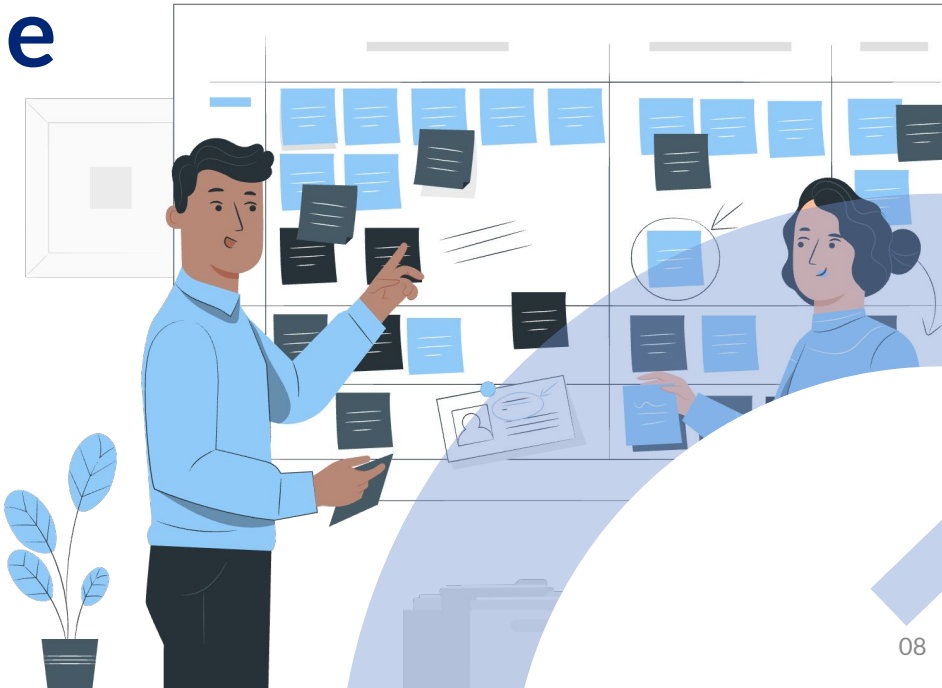

The P × I model and its structural blind spot

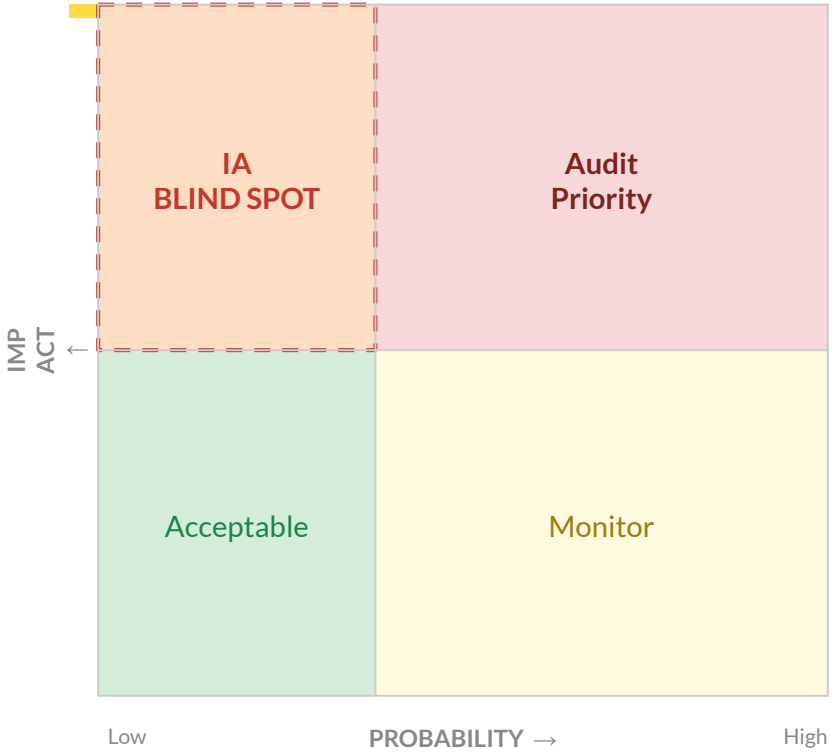

The flaw:

The heat map correctly colors the blind spot orange.

But organizations still ignore it, because the 'Low Probability' label gives permission to put it off, push it down the list, or formally accept it.

No board ever asks: 'What if our probability estimate is wrong?'

The color is correct. The behavior is wrong.

OpenQ

P × I was designed for a stable world.

→ Labels stick: once a risk is rated 'low,' it stays low year after year. Nothing forces a fresh look.

→ Set once a year: probability is judged in December for all of next year. The world moves faster than the calendar.

→ Judged by label, not by cost: audit committees accept a risk by its label ('low probability'). They do not ask what happens if that label is wrong.

EY 2026 calls this the NAVI problem: risks that are Nonlinear, Accelerated, Volatile, and Interconnected.

The world we are actually operating in

N — Nonlinear

Risks trigger sudden tipping points with no warning. Small events cascade into systemic shocks.

A — Accelerated

The speed of threat development now outpaces the speed of organizational response and audit cycles.

V — Volatile

Direction of risk shifts rapidly and unpredictably. Last year's heat map is already out of date.

I — Interconnected

A single trigger sets off several risk categories at once: geopolitical, operational, cyber.

EY 2026 Audit Committee Priorities Report · NAVI risk mentions in public disclosures up 50% since 2020

You don’t wear a seatbelt because crashes are likely.

You wear it because the consequence of not wearing it is unacceptable.

IA does not audit risks because they are likely.
IA audits risks because the organization cannot survive if they happen.

OpenQ

A different question changes everything

Old question:

"How likely is this risk?"

New question:

"Can we survive if this happens?"

Section 2 shows what happens when we apply this question to real research data and Thai companies.

OpenQ

Finding the single point: trace the dependencies

A single point of failure is not always technical. It can be one vendor, one cable, one undocumented script, or one person who knows how it works.

MAP

Critical services

List the services the business cannot run without for even an hour.

Business-critical services

TRACE

Follow dependencies

Map each service down to vendors, infrastructure, and shared components.

Vendors, infra, people

FIND

Spot single nodes

Where do many services converge on one vendor, one cable, one component?

Convergence points

TEST

Assume it fails

For each single point, run the failure scenario. Can you recover without it?

Failure scenario test

Low probability is not the question.

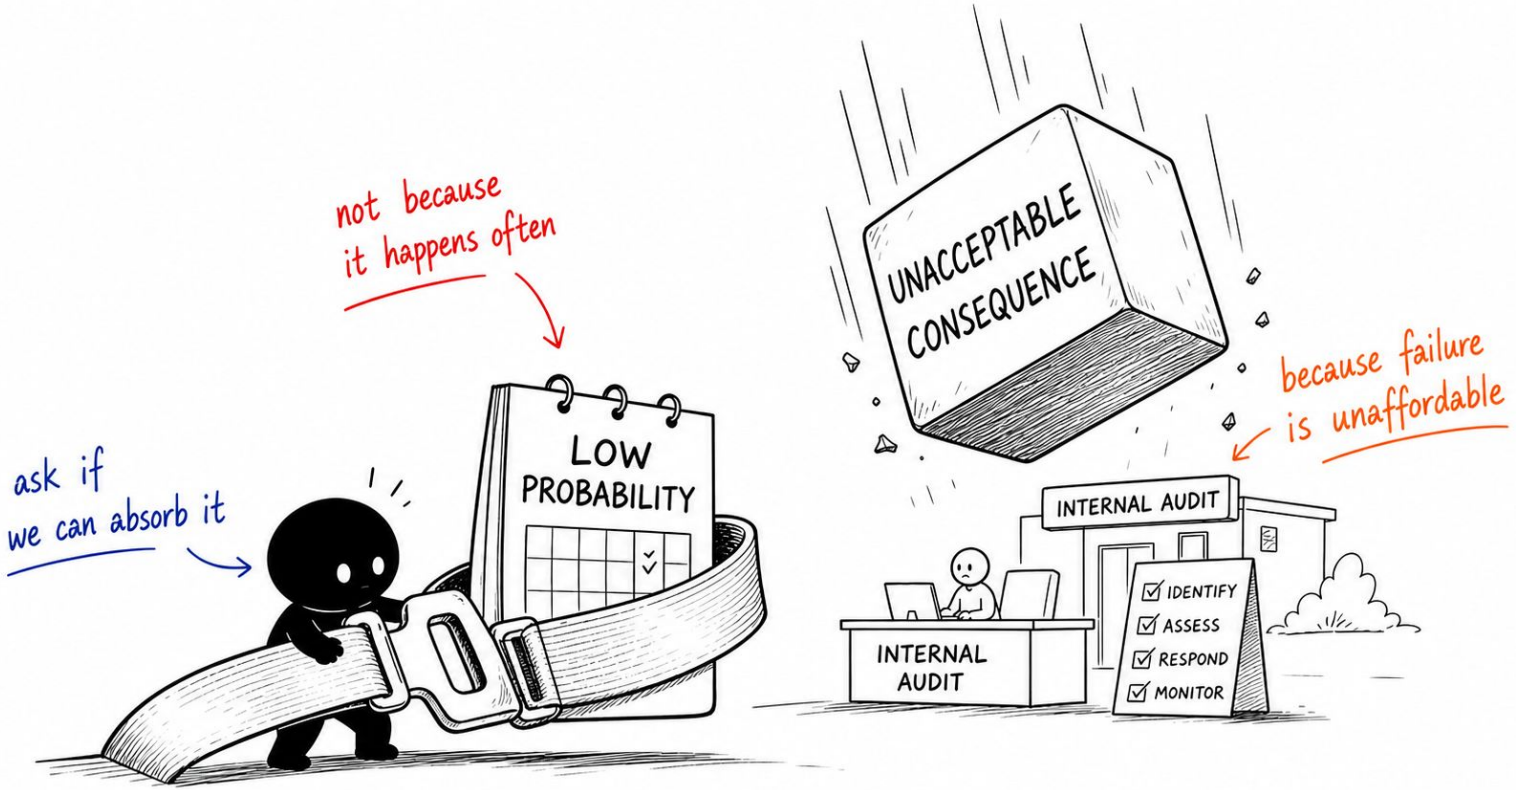

Q: Is the annual risk assessment still enough?

- No. Annual is the floor, not the plan.
- The world changes faster than December planning can track.
- The P x I model has a structural blind spot: low probability gives permission to ignore.
- Section 2 shows real cases where the blind spot caused the damage.

OpenQ

Hot topics and risks 2026

02

Big4 research findings + Thai SET50 company evidence

Thailand's financial sector: structural change in 24 months

2024–2025: more regulatory and market change than the previous five years combined

MARKET STRUCTURE CHANGES

Early 2024
Global crypto exchange enters via JV with Thai-listed energy company. A new competitor in the licensed market.

Early 2024
International crypto platform suspends all Thai trading. Customers move to local licensed operators.

2024
Commercial bank acquires and rebrands a licensed crypto exchange. A bank-owned crypto platform is created.

April 2025
Major global exchange enters via an existing Thai licensee. It becomes the ninth licensed operator in Thailand.

May 2025
Government moves to block 5 unlicensed global platforms. A grace period is given for Thai users to retrieve assets.

REGULATORY FRAMEWORK CHANGES

Aug 2024
New governance framework: stricter director qualifications and risk management structure required for all operators

Jan 2025
First-ever approval for mutual and private funds to invest in digital assets. An entirely new product scope.

March 2025
USDT and USDC approved as eligible trading assets. Stablecoins enter the regulated product list.

April 2025
Rules reach offshore: platforms targeting Thai users now need a Thai license or face criminal penalties

July 2025
Securities companies allowed to offer investment token trading. The boundary between sectors is now blurred.

The audit scope that did not exist in January 2024

If your risk register and audit plan were written at the start of 2024, every item below was missing from scope.

New entities to audit

- Operators that entered the market after the plan was written
- Commercial bank-owned crypto platform: cross-entity risk controls
- Operators that exited: customer asset protection and wind-down
- Virtual banks begin June 2026: no regulatory precedent in Thailand

New regulations to test

- 6+ SEC notifications in 18 months: each one needs compliance testing
- First amendment to Digital Asset Decree since 2018 published Aug 2024
- Rules now reach offshore: do your entities need Thai licenses?
- Funds in digital assets: new portfolio risk, new controls required

New models, no playbook

- Bank-owned crypto: securities + digital asset boundaries blur
- Virtual banks: digital-only, no branch BCP, so a different resilience model
- Securities firms in token trading: conflict of interest, scope overlap
- BOT stablecoin sandbox: new payment rails below IA radar

An audit plan written in January 2024 would have missed every item above.
Not because IA failed, but because the industry moved faster than any annual cycle can track.

What the Big4 agree on for IA in 2026

DELOITTE 2026

"Most likely AND most severe": two separate audit tracks

Resilience is not just regulatory compliance. It is an embedded organizational capability. IA must test both categories.

EY 2026

The NAVI world: Nonlinear, Accelerated, Volatile, Interconnected

Only 12% of non-financial Fortune 500 companies have a Chief Risk Officer. Boards demand "war gaming" and what-if scenario analysis.

KPMG 2026

Polycrisis: the combined effect is bigger than the sum of the parts

BCPs must be tested against "natural disasters, cyberattacks, and supply chain disruptions." KPMG names these explicitly.

PwC 2026

"Businesses cannot plan for every scenario"

Resilience = building capacity to absorb shocks and adapt quickly. The goal shifts from comprehensive planning to survivability.

IIA Risk in Focus 2026 · 4,073 respondents · 131 countries · Geopolitical risk: largest single-year jump (+10pts) of any risk category

Deloitte 2026 Hot Topics for Operational and IT Internal Audit

> **IA should test whether the business can withstand the most likely AND most severe disruptions
> Two separate categories, not one.**

The key word is AND. Not 'the most likely.' Not 'the most severe.' Both, in parallel.

Deloitte Hot Topic #6: Resilient Business. IA must review how complete the scenario planning is, test cyber and climate readiness, and check that resilience is built into day-to-day operations, not treated as a once-a-year exercise.

Source: Deloitte 2026 Hot Topics for Operational and IT Internal Audit (UK, published Jan 2026)

EY 2026 Audit Committee Priorities Report

12%

of non-financial Fortune 500 companies have a Chief Risk Officer

The EY NAVI framework

- **N** Nonlinear
- **A** Accelerated
- **V** Volatile
- **I** Interconnected

EY recommends boards use scenario analysis and war gaming, asking 'what if?' then 'what can?' to prepare for connected risks that could combine into a big shock.

Source: EY 2026 Audit Committee Priorities Report

Three findings that validate the session's core argument

KPMG 2026

> “BCPs must be tested against realistic scenarios such as natural disasters, cyberattacks, and supply chain disruptions.”

KPMG recommends IA check whether BCPs are up to date and test that tabletop exercises cover all three scenario types, not just cyber.

PwC 2026

> “Businesses will not be able to plan for every scenario. Resilience means building the capacity to absorb shocks and adapt quickly to new realities.”

PwC Risk Agenda 2026. The shift from comprehensive planning to resilience capacity is the session's core methodology argument.

IIA RISK IN FOCUS 2026

+10 pts

Geopolitical risk: the single largest year-on-year jump of any category. Based on 4,073 Chief Audit Executive responses across 131 countries.

Across Big4 and IIA research: risk assessment cycles are not keeping pace with the speed at which the environment is changing.

Six threats beyond geopolitical and natural disaster

Each is correctly colored orange or red on the heat map. None appear in any Thai SET50 56-1 disclosure.

01

Quantum / Q-Day

Q-Day estimated 2030: all current encryption at risk

02

Critical infra cyber attack

Lloyd's $3.5T scenario: payment systems taken down at once

03

Pandemic / biological

H5N1 active 2025: global readiness rated 'increasingly fragile'

04

Systemic AI model failure

Correlated failure across the sector: ECB May 2026 stability warning

05

AI deepfake bank run

SVB fell in 48h via social media: AI moves faster than checks can

06

Submarine cable failure

One tanker cut 5 cables in 2025: Thailand's payments run on these

OpenQ

2030.

The year quantum computers are expected to break current RSA and elliptic-curve cryptography, the encryption that protects every Thai financial transaction.

“Harvest Now, Decrypt Later.”
Attackers are already collecting encrypted data today to decrypt once Q-Day arrives.

No Thai IA function currently has Quantum / Q-Day in its audit plan. It has a hard deadline.

Sources: Deloitte 2026, Forrester 2026, NIST IR 8547, Google AI Q-Day forecast

OpenQ

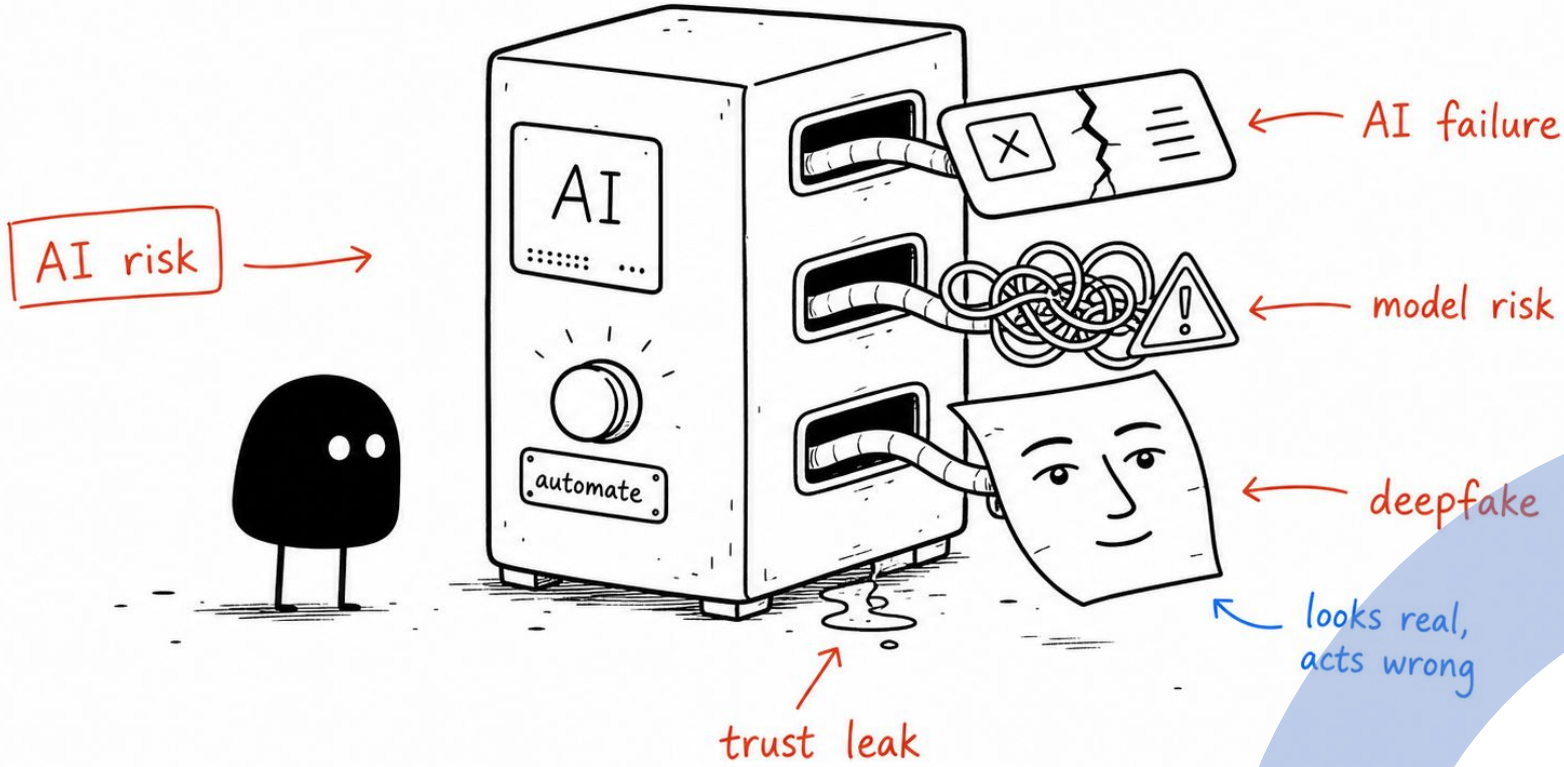
AI risk
AI
automate
AI failure
model risk
deepfake
looks real,
acts wrong
trust leak

Systemic AI model failure

Several Thai financial institutions use the same AI model for credit scoring or fraud detection and fail at the same time. This correlated failure is invisible to any single firm's risk assessment. ECB Financial Stability Review, May 2026: frontier AI is an explicit risk to financial stability.

AND

AI deepfake-driven bank run

Silicon Valley Bank collapsed in 48 hours through social media coordination in 2023, without any AI help. A deepfake of a Thai central bank governor or major bank CEO announcing insolvency, sent out at 3am, could trigger the same run before anyone could check it.

Sources: ECB Financial Stability Review May 2026 · WEF Global Risks Report 2025 (#1 medium-term risk: misinformation)

OpenQ

$3.5T

Estimated global economic loss from a coordinated cyber attack on financial infrastructure. Lloyd's of London systemic risk scenario.

→ In 2025, a single tanker linked to sanctions-evading networks cut 5 submarine cables. Estlink 2, the electricity cable between Finland and Estonia, was offline for 7+ months.

→ Cloudflare tracked 180+ internet disruptions globally in 2025. Thailand's entire digital payment and interbank settlement system depends on undersea cable corridors.

Sources: Lloyd's Systemic Risk Scenario · Gallagher Report 2026 · Cloudflare Q4 2025

The full risk landscape: where are the 8 blind spot threats?

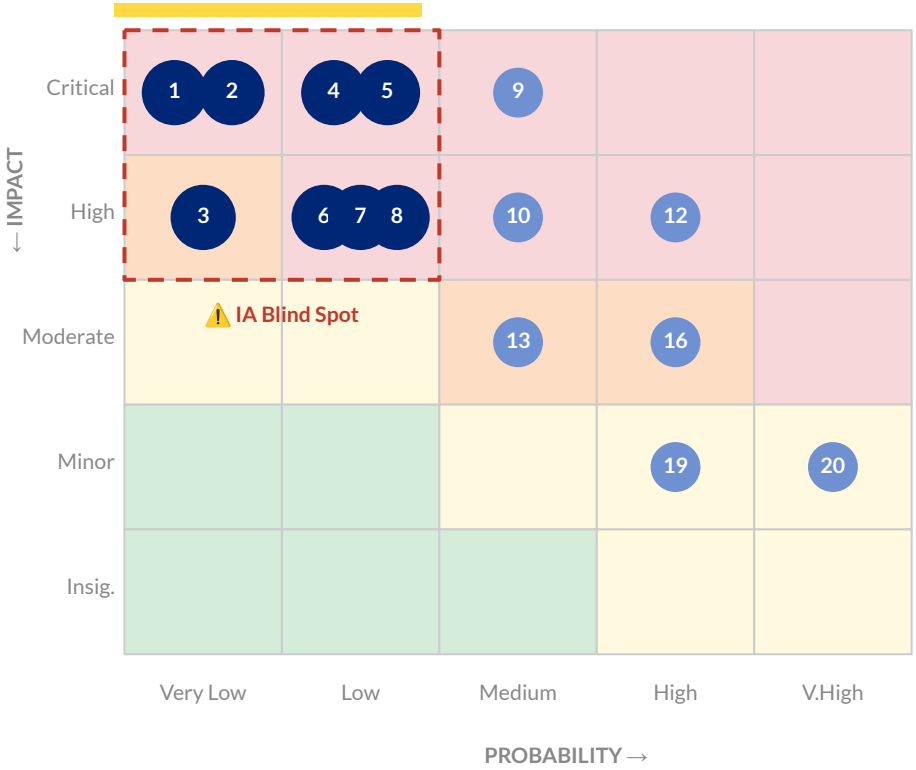

IA blind spot risks (numbered)

1. Quantum / Q-Day
2. Financial system collapse
3. Major natural disaster
4. Pandemic / biological outbreak
5. Critical infrastructure attack
6. Systemic AI model failure
7. Submarine cable failure
8. AI deepfake bank run

Lighter dots = risks already on most audit plans

What Thai SET50 companies disclose in their 56-1 One Reports

Risk category	SCBX	KBANK	PTTGC	CPALL
Cybersecurity	✓	✓	✓	✓
Geopolitical / trade risk	✓	✓	✓	✓
Climate / ESG risk	✓	✓	✓	✓
Regulatory compliance	✓	✓	✓	✓
BCM / operational safety	✓	✓	✓	✓

All four companies disclose these categories. This is the regulatory minimum required by SEC Thailand and the BOT for listed financial entities.

OpenQ

What about these risks?

Check these against your own audit plan, right now.

☐ **Seismic / natural disaster**	Is there a scenario-based BCM audit in your plan that tests an event that has never happened before?
☐ **Quantum / Q-Day**	Does your audit plan include a review of cryptographic risk and readiness for when current encryption breaks?
☐ **Systemic AI model failure**	If your institution and five peers all use the same AI model and it fails at the same time, is that in your scope?
☐ **AI deepfake / social media run**	Is there any audit procedure that would detect or respond to a social-media-driven confidence crisis?
☐ **Submarine cable / internet dependency**	Has IA validated your BCP against a scenario where Thailand's internet infrastructure is disrupted for 72 hours?

If any answer is 'not sure' or 'no,' you are not alone. Section 3 shows exactly what to do about it.

The chain that creates the blind spot

01

Regulators require disclosure of X

Companies include X in their 56-1 One Report

02

Management builds risk register

Risk register follows what is in the disclosure

03

IA plans audits against the register

Audit plan covers what management registered

04

What is not required is not registered

What is not registered is not audited

This chain is not a failure of bad intent. It is a structural limitation of disclosure-driven risk management. The role of IA is to break this chain by independently scanning the risk landscape and bringing what is missing into scope.

Seismic risk. Quantum threat. AI systemic failure. None required by regulators. None in the register. None in the audit plan.

OpenQ

The blind spot sits behind the chain.

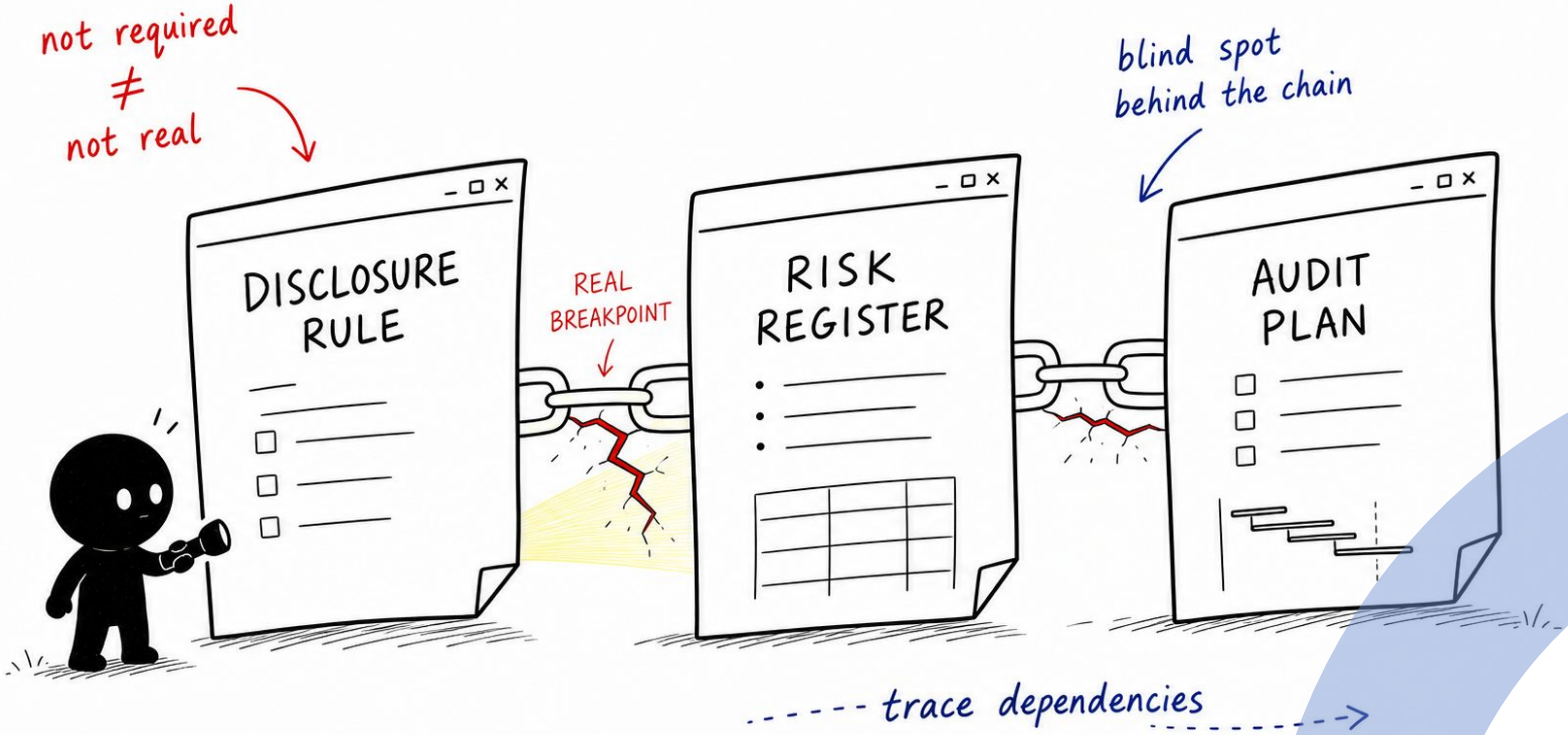

OpenQ

Q: Does this only apply to financial services?

- No. The pattern is the same in every industry.
- Finance, insurance, manufacturing, aviation, public sector.
- The common factor is not the sector. It is hidden dependency.
- One vendor, one cable, one model failing, and it reaches everywhere.

This pattern happens in every industry

Finance and insurance

- New products appear faster than the audit plan
- Crypto, virtual banks, AI credit scoring
- Regulator issues guidance mid-year
- Your scope changes before the plan does

Manufacturing, aviation, public sector

- Single supplier fails, whole line stops
- Aviation: one vendor update, global groundings
- Public sector: disruption labeled low risk
- University: AI changes research integrity risk

The risk is not the industry. It is any dependency you cannot see until it fails.

OpenQ

03

The operational blind spots

Where incidents really come from, and why our tools miss them

30 days.

The usual reporting cadence for a key risk indicator. By the time a monthly KRI turns red, the incident it was built to warn about has often already happened.

A KRI that counts incidents after they occur is not an early warning. It is a historical record with a reporting delay built in.

The question for IA is simple: does this indicator move before the risk happens, or only after?

OpenQ

Most KRIs are lagging. The useful ones are leading.

Lagging KRI

Counts the past

Counts incidents that already happened

Reported monthly or quarterly

Threshold reviewed once a year

Shows the damage, not the warning

→

Leading KRI

Signals early

✓ Tracks near-misses and failed attempts

✓ Measures velocity: how fast it is changing

✓ Alerts on a threshold, not on a calendar

✓ Moves before the incident, not after

OpenQ

Every AI agent adopted this year is a new vendor dependency

How AI tools enter the org today

No register

Procured by a team, not by IT or risk

No record of which process now depends on it

Cancelling the subscription is assumed to be simple

No owner accountable if the tool disappears

→

What IA should be asking instead

An exit plan

✓ Which workflows require this agent to run?

✓ n8n surged then faded. The workflows on it did not.

✓ If cancelled next year, what breaks?

✓ Was the manual fallback deleted when this shipped?

OpenQ

The AI exit readiness checklist

INVENTORY

Is there a register of every AI tool in production?

Not just chat tools. Includes embedded agents, automation platforms, and AI features inside SaaS products already approved for other reasons.

OWNERSHIP

Does each tool have a named business owner, not just an IT ticket?

If the answer is the procurement record only, no one is accountable for what happens when the subscription ends.

FALLBACK

Is there a manual or alternate path if the tool is cancelled tomorrow?

n8n adoption surged, then slowed sharply as teams hit reliability and support gaps. Workflows built on it had no fallback when it did.

EXIT COST

Has anyone priced what unwinding this tool actually costs?

Cost is not just the subscription. It is rebuilding the process, retraining staff, and the gap while that happens.

The IA question: if every AI subscription were cancelled tomorrow, which processes would simply stop working?

OpenQ

The change looked isolated. The blast radius was not.

What the change ticket claims

Looks isolated

Affects one system, low risk

Reviewed by the team that owns it

Rollback is assumed to be simple

Downstream dependencies not listed

→

What the blast radius really is

Maps the spread

✓ Every system downstream of the one changed

✓ Shared services, payment rails, customer apps

✓ Before a change: what could this take down?

✓ After an incident: what did it actually reach?

OpenQ

Fixing the symptom is why the incident comes back

Symptom fix

Stops the bleeding

Restarts the server, closes the ticket

Blames the person who made the change

Addresses what is easy to see

Same incident returns next quarter

→

Real root cause

Prevents recurrence

✓ Asks why five times, not once

✓ Why was there no second review?

✓ Fixes the process, not the person

✓ CrowdStrike: the process, not the file

Four moves on the operational blind spots, starting today

None of this needs budget or new systems. It needs IA to ask a different question about work the organization already does every day.

KRI

Make one leading

Take one lagging KRI. Add a near-miss or velocity measure that moves first.

Leading indicator

CHANGE

Audit a routine one

Pick one normal change. Was impact assessed? Was there a second review?

Change as risk event

SPOF

Map one service

Trace one business-critical service to its single points of failure.

Dependency map

RCA

Reopen one incident

Take a closed incident. Ask why five times. Cause or symptom?

Real root cause

Q: How do we keep up when risks change faster than our audit cycle?

- Stop waiting for the annual register. Add a quarterly horizon scan.
- Use leading KRIs, not just lagging ones.
- Map single points of failure before the incident, not after.
- Section 4 gives you the operating model and the tools to do this.

OpenQ

04

How IA responds

The methodology shift: from probability to survivability

It happened. The board turns to you.

Will you call it bad luck?

It was never bad luck.

The signal was inside your scope.
No one was assigned to read it.

OpenQ

300

“On average, around 300 warning signals precede one serious failure (Bird, 1969). They are almost always visible in advance.”

Makkasan, Bangkok, 16 May 2026. A freight train hit a stranded bus:

8 killed, 30+ injured. Warned in print for years.

The same failure killed 18 people at a Chachoengsao crossing in 2020.

The signal did not whisper. It had already screamed, once.

Sources: F. Bird (1969) industrial loss study; Makkasan collision 16 May 2026 and Chachoengsao collision Oct 2020 (Bangkok Post)

In 2024, the profession added one word to its own standard.

Foresight.

Here is how internal audit delivers it.

OpenQ

Stop asking how likely. Start asking can we survive.

Traditional IA

Lagging indicator

- ✗ Audits what already happened
- ✗ Plans against last year's register
- ✗ Reports findings after the fact
- ✗ Asks: how likely was this?

→

Adaptive IA

Leading indicator

- ✓ Audits what could happen
- ✓ Plans against what the landscape actually contains
- ✓ Reports emerging exposure before it happens
- ✓ Asks: can we survive if this happens?

OpenQ

Stop asking how likely. Start asking can we survive.

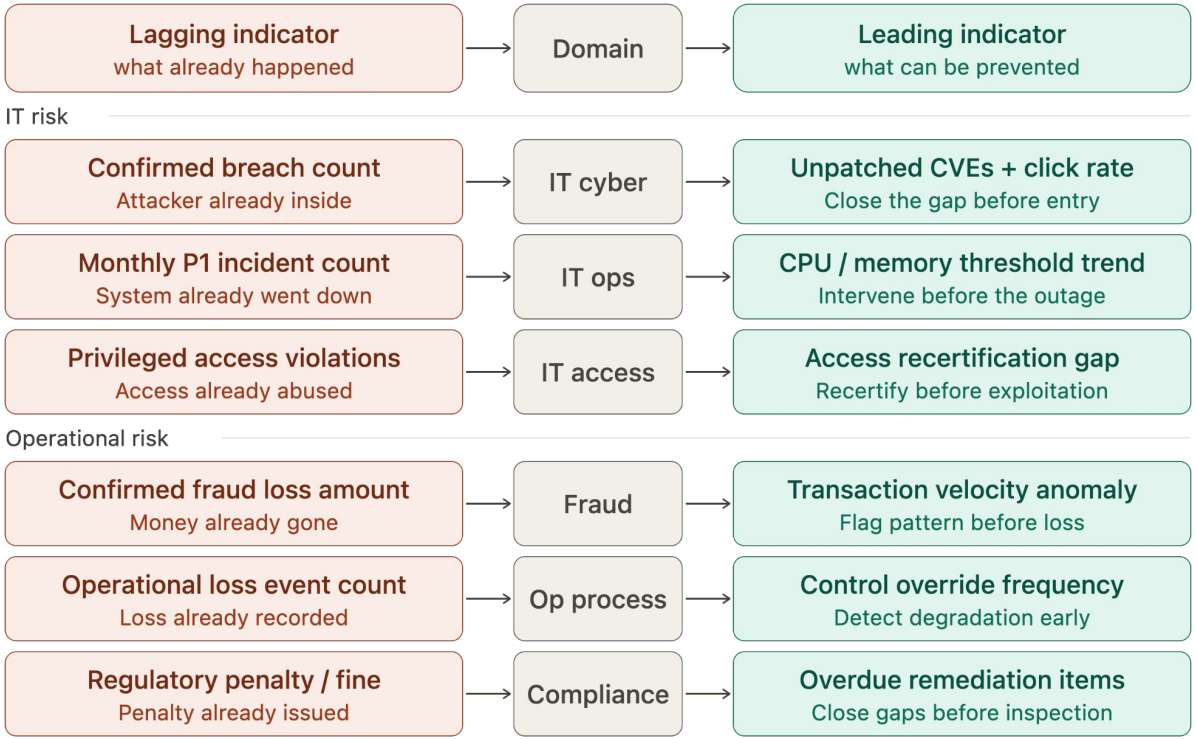

The adaptive IA operating model

This is not a theory. Each step is a decision you can make this week without budget approval.

SEE

Continuous sensing

Continuous monitoring and leading KRIs, scanning BOT, WEF, CISA and news against the register.

Continuous monitoring + KRIs

→

RESPOND

Re-plan and execute

Re-prioritise the backlog as signals move. Test scenarios in short sprints. (Agile IA)

Agile IA: re-prioritisable backlog

→

OWN

Assign and assure

Name who owns each signal (Three Lines). Report the coverage gap to the board quarterly.

Three Lines: named accountability

The unacceptable floor: three questions to ask your board now

Some risks are unacceptable regardless of their probability score.

Bring these three questions to your next audit committee meeting:

→ What scenario, if it happened, would end this organization?

→ What risk, if undetected, would make every other finding irrelevant?

→ What has never happened here, but would matter most if it did?

Example floor list, Thai financial institution:

- Seismic / physical infrastructure failure
- Critical payment system outage (>4 hours)
- Quantum breach of encryption in production
- Social-media-driven confidence run
- BOT/SEC emergency regulatory action

The floor list is not part of the heat map. It sits above it.
These risks are always monitored regardless of probability rating.

Scenario-based execution: the survivability test

Old audit step:

"Is this control designed and documented?"

New audit step:

"If this scenario happened at 9am on a Monday, would our controls detect it in time to respond?"

Bow-tie analysis

Map causes AND consequences, not just likelihood. Test controls on both sides of the event: prevention AND recovery.

BCM validation

BCM drills become audit scope. IA tests whether the drill covered the right scenario, not just whether a drill was documented.

Sprint execution

2-4 week focused engagements per scenario. Replaces 3-month engagements that test too many controls too superficially.

The coverage gap: what to show the board

What IA currently covers	What the risk landscape contains
✓ Cybersecurity controls	✗ Seismic / natural disaster (scenario-tested)
✓ Regulatory compliance	✗ Quantum / Q-Day encryption threat
✓ Credit and market risk	✗ Systemic AI model failure
✓ Supply chain and third-party risk	✗ AI deepfake bank run scenario
✓ BCM (documented, not scenario-tested)	✗ Submarine cable / internet dependency
✓ Climate / ESG disclosures	✗ BCM scenario that has never happened

Present this to the audit committee. Ask: “Which of these should we be covering?”
That question changes the conversation.

OpenQ

5 things IA does differently, starting Monday

No budget required. Each action has a clear output you can show your committee.

01 **Define the unacceptable floor list**

Agree 5-7 scenarios that IA monitors regardless of probability. Present it to the audit committee this quarter.

02 **Ask the survivability question at the next AC meeting**

“What risk, currently rated low probability, would end us if it happened?” One question. Let it breathe.

03 **Add a quarterly horizon scan to the audit cycle**

Read WEF, BOT guidance, and one Big4 hot topics report each quarter. Map gaps to your register. 2 hours per quarter.

04 **Rebuild two audit steps as scenario-based**

In the next engagement, add: “If scenario X happened, would these controls have detected it in time?” Start with BCM.

05 **Show the coverage map to the board**

Two columns: what IA covers vs what the risk landscape contains. The gap is the conversation. It should happen every quarter.

OpenQ

Q: What skills does IA need to handle these new risks?

- Three layers: audit judgment (unchanged), data and AI literacy, scenario thinking.
- You do not need to be a data scientist. You need to spot a wrong answer.
- When you lack depth, borrow it. Co-sourcing is in the 2024 IIA Standards.
- Section 5 shows where most Thai IA functions are today and the path forward.

OpenQ

Three skills IA needs right now

You do not need to become a data scientist. You need enough to ask the right question.

- **Audit judgment (unchanged)** -- Your professional opinion is still what the board pays for. AI does not sign the assurance opinion. You do.
- **Data and AI literacy** -- Read a model's output critically. Spot a wrong answer. Know what the tool cannot do. Use it without depending on it.
- **Scenario thinking** -- Design the unacceptable floor list. Run a tabletop logic test. Ask: if this happened at 9am Monday, what breaks first?
- **When you lack depth: borrow it**. Co-source a specialist for one engagement. Ask the right questions. Validate the answers. The judgment is still yours.
- The 2024 IIA Standards (Standard 9.5) require IA to coordinate with other assur providers. Co-sourcing is not a weakness. It is the standard.

Three categories of IA monitoring tools

OpenQ

What you want to monitor	Concept	Tools / examples	Right for
External risk signals New threats and events not yet in your register	**AI sensing** AI scans many sources and flags what your register is missing	**Level 1: Claude + prompt** Level 2: Python + Claude API Level 3: MCP platform	Every IA team. Zero cost.
Internal process deviations Unapproved steps and control bypasses inside your systems	**Process mining** Reads ERP event logs and shows where the process went wrong	**ACL / IDEA / Power BI** Celonis / SAP Signavio IBM Process Mining	Mid to large orgs with ERP and a data team
Scenario simulation Run a what-if before the event happens	**Swarm intelligence** AI agents simulate how people and systems react to an event	**MiroFish (open source)** Research-grade. Use for scenario design, not production audit	Scenario planning teams. Complements the other two.

Start with row 1 today. Add process mining when your data infrastructure is ready.

OpenQ

How IA responds, part 2

04

Leveraging agentic AI: the sensing engine that frees IA to focus on judgment

Which tool fits your team right now?

OpenQ

Team size	Budget	Start with this	Move to this when ready
Any size	Zero	**Claude + risk sensing prompt** Upload register. Run prompt. 2 hours.	Python + Claude API for weekly automation. No IT approval needed.
5 to 15 people	Low	**ACL / IDEA / Power BI** Connect to ERP extracts. Build control monitoring dashboards.	Add Claude API sensing layer on top. Two streams running in parallel.
Large bank or enterprise	Medium to high	**Celonis / SAP Signavio** Full ERP event-log mining. Needs IT project team. 6 to 18 months to deploy.	Add MiroFish-style simulation for scenario stress testing alongside live monitoring.

The most common mistake: waiting for budget to start. Row 1 costs nothing. Start there today.

OpenQ

The information bottleneck IA cannot manually solve

To stay ahead of emerging risks, IA would need to monitor:

WEF Global Risks Report	Annual, but updated quarterly online
Big4 hot topics (4 firms)	Multiple reports per firm, per year
BOT and ก.ล.ต. circulars	Unpredictable: issued as needed
CISA and MITRE threat alerts	Weekly, sometimes daily
USGS / GDACS physical risk feeds	Real-time
Industry news and OSINT	Continuous

Manual scanning across all these sources:

6-10 hours per week

per analyst

That time is not being spent on judgment, analysis, or assurance.

AI removes the information burden. Humans retain the assurance judgment.

The AI risk sensing system: how it works

DATA SOURCES

Thai regulatory
BOT · ก.ล.ต. · SEC Thailand

Global risk intel
WEF · Big4 · IIA · CISA

News & OSINT
Reuters · Thai media

Physical risk
USGS · GDACS · weather

Org context
Risk register · audit plan

INTELLIGENCE ENGINE

⏰ Weekly schedule trigger (Monday 6am)

Claude AI
Signal extraction · Gap analysis · Risk scoring

MCP tools available:

web_search

web_fetch

Gmail MCP

risk_register_mcp

regulatory_mcp

WEEKLY OUTPUTS

Weekly risk brief
Evidence-based narrative + sources

Watchlist delta
New risks in · resolved risks out

Audit triggers
Plan adjustment required: act now

Coverage gap map
Register vs real landscape

AI senses. IA judges.

OpenQ

What the weekly brief actually looks like

WEEKLY RISK INTELLIGENCE BRIEF

AUTO

Week of [date] · AI-assisted · Reviewed by CAE before distribution

TOP NEW SIGNALS THIS WEEK

QUANTUM	NIST published updated PQC migration timeline. Financial institutions advised to audit cryptographic dependencies by Q1 2027.	NOT IN CURRENT AUDIT PLAN
GEOPOLITICAL	ASEAN energy market disruption: oil futures +12% week-on-week. BOT foreign reserve guidance updated.	Covered: monitor
CYBER	CISA advisory: state-sponsored actors pre-positioning in SWIFT-connected financial infrastructure.	AUDIT TRIGGER: see below

Watchlist: 12 active risks · 2 new this week · 1 resolved · Audit triggers: 1 item requires plan adjustment this quarter

Generated by Claude AI + MCP tools | Human review by CAE before any action | Distributed: CAE, Deputy CAE, Audit Committee Chair

This is a fictional illustrative example. The actual content is generated fresh each week from live data sources specific to your organization.

OpenQ

The independence line: what IA can and cannot

IA can do this	IA cannot do this
• Monitor risks in real time • Flag emerging signals to management • Advise on control design (without deciding) • Attend risk discussions as an observer • Run scenario tests and share findings • Report the coverage gap to the board	• Choose which risks the business accepts • Own a control or sign off on its design • Make operational decisions for the business • Manage the risk register on behalf of management • Approve a risk response

The line is simple:
IA flags and assures. Management decides and owns.
Cross that line and you cannot audit it. (IIA Standard 9.5 + Three Lines Model)

OpenQ

Three maturity levels: where can you start?

Today: no build required

Assisted intelligence gathering

Analyst + Claude ·
2-3 hrs / week

- Upload your risk register + relevant reports into Claude
- Ask Claude to identify gaps and new signals
- Review output and distribute manually
- No API keys, no code, no scheduling

1-2 weeks to build

Script-assisted automation

Python + Claude API ·
15-30 min / week oversight

- Python script reads Excel register + fetches free APIs
- Calls Claude API with full context weekly
- Outputs structured brief + emails via Gmail
- Runs via Windows Task Scheduler or cron

4-8 weeks to build

Full MCP intelligence platform

MCP servers + n8n + Claude ·
5 min / week oversight

- Custom MCP servers for each data source
- Automated weekly brief with full source coverage
- Live coverage gap dashboard for audit committee
- Multi-organization configuration possible

Start at Level 1 this week. The most common blocker at Level 2 is getting IT to approve the API key. Plan for that early.

OpenQ

The right division of labour between AI and IA

AI does

→ Scans hundreds of sources in minutes

→ Extracts risk signals from unstructured text

→ Compares signals against the risk register

→ Scores and prioritises by severity and velocity

→ Tracks delta from previous week

→ Formats the weekly brief

→ Distributes via email or Slack automatically

IA does

✓ Reviews AI output with professional scepticism

✓ Applies context that AI cannot know

✓ Makes the judgment call: audit or accept?

✓ Asks the survivability question to the board

✓ Signs the assurance opinion

✓ Builds the relationship with leadership

✓ Takes accountability for what is covered and what is not

AI removes the information burden. IA keeps the assurance responsibility. The IA function is not replaced. It is elevated.

Getting started: a prompt you can use today

Level 1 requires no code, no API key, no budget. Just Claude and your risk register.

Prompt template: Level 1 risk sensing (copy, adapt, use today)

You are a risk intelligence analyst for [ORGANIZATION NAME], a Thai [sector] regulated by BOT / ก.ล.ต.

I am attaching our current risk register [attach file]. Today's date is [date].

Your task:
1. Identify any risk signals in the past 30 days that are NOT covered in our current register
2. Flag which of these could be classified as Low Probability, High Impact
3. For each gap found, state: what the risk is, what the evidence is, and why it matters for IA
4. Recommend whether any should be added to our Q3/Q4 audit plan

Focus especially on: seismic / physical events, emerging technology threats, systemic risks affecting the Thai financial sector, and any BOT or ก.ล.ต. guidance issued recently.

Step 1

Upload your risk register to Claude. Paste this prompt. Read the gap analysis.

Step 2

Share the output with your audit committee. Ask: which gap should we close first?

OpenQ

Measuring maturity and roadmap

05

Where your IA function is today, and the path to where it needs to be

OpenQ

Internal audit maturity model: COBIT-style 5-level scale

DIMENSION	L1 Initial	L2 Managed	L3 Established	L4 Predictable	L5 Optimizing
		Most Thai IA today	12-month target	AI-enabled	
D1 Risk sensing Horizon scanning	Register only: no independent scan	Annual external review: mainly internal sources	**Quarterly horizon scan + gap vs register**	Weekly AI intelligence brief: multi-source	Real-time predictive risk identification
D2 Audit planning Prioritization	Ad hoc: reacts to incidents	Annual P×I heat map: fixed for 12 months	**Dynamic plan + quarterly trigger reviews + floor list**	AI signals adjust plan continuously	Self-adjusting: board validates quarterly
D3 Execution Methodology	Control existence: is it documented?	Control operation: is it working?	**Scenario-based: would it survive X?**	Resilience testing + bow-tie analysis	Continuous real-time monitoring
D4 Assurance output Communication	Findings report: backward-looking	Annual regulatory assurance opinion	**Quarterly coverage gap report to board**	Survivability assurance: not just compliance	Live board dashboard: continuous
D5 Technology & AI Enablement	Word / Excel: no analytics	Audit management system: manual tools	**AI-assisted research and documentation**	AI-powered sensing platform: MCP-connected	Agentic AI: autonomous procedures

Click any cell to read the full descriptor. See the interactive version in the session resources.

Where most Thai IA functions sit, and why

Why most Thai IA functions are at Level 2:

→ The annual P×I heat map is the main way risk is prioritised, and it anchors every year

→ Regulatory compliance (BOT, SEC) sets the minimum, and in practice it becomes the ceiling

→ Low probability items are accepted or put off without question: the heat map gives permission

→ IA plans against management's register, and the register follows what regulators require

The gap from L2 to L3 is a methodology decision, not a budget decision. It requires changing what IA asks, not what IA spends.

OpenQ

Your 12-month action list: from where you are to where you need to be

This quarter
Immediate: zero budget

1. **Define unacceptable floor list**
   5-7 scenarios always monitored regardless of probability
2. **Ask survivability question at AC**
   "What risk would end us if it happened?"
3. **Start Level 1 risk sensing**
   Upload register to Claude. Read the gap. 2 hours.

3 months
Q4 2026: process change

1. **Scenario-based steps in 1 engagement**
   Add 'would this control survive X?' to next audit program
2. **Quarterly horizon scan established**
   WEF + BOT + one Big4 report per quarter mapped to register
3. **Coverage gap report to AC**
   Two columns: what IA covers vs what landscape contains

6 months
Q1 2027: tooling

1. **Level 2 AI sensing running**
   Python + Claude API: automated weekly brief
2. **Dynamic plan with trigger reviews**
   Quarterly plan refresh based on AI signals
3. **BCM scenario validation in scope**
   At least one scenario that has never happened before

12 months
Q2 2027: IA at L3

1. **Coverage map to board quarterly**
   Standard agenda item: board sees the full risk landscape
2. **IA function rated L3 across all 5 dimensions**
   Risk sensing · Planning · Execution · Assurance · Technology
3. **Build the case for Level 3**
   AI-powered sensing platform: investment proposal ready

OpenQ

Q: How do we justify budget for risks that may never happen?

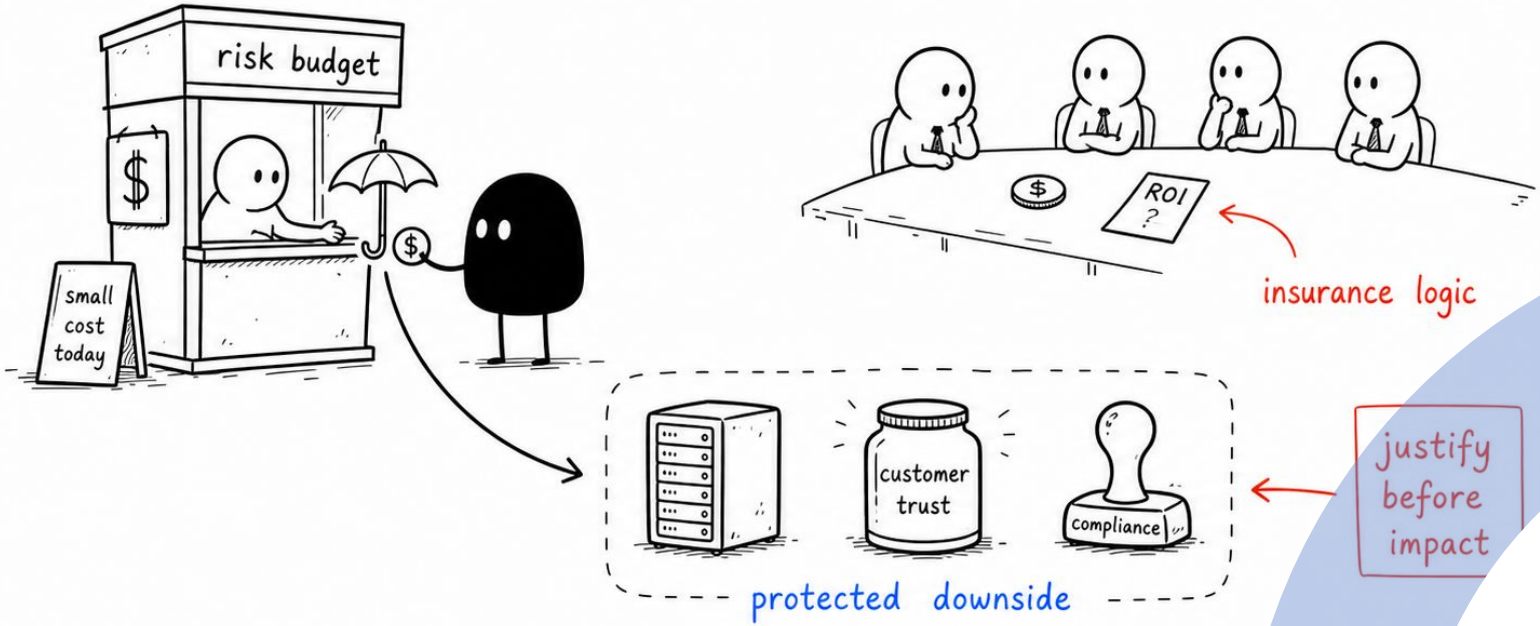

OpenQ

Q: How do we justify budget for risks that may never happen?

- Do not sell probability. Sell survival.
- The board will not buy 0.1% chance. It will buy: we cannot afford to be wrong.
- Show the coverage gap: what IA covers vs what the landscape contains.
- That two-column slide changes the conversation in every audit committee.

Three things. This quarter. Zero budget.

01 This week: 2 hours

Define your unacceptable floor list

Write 5-7 scenarios your organization cannot survive. No probability rating needed. Output: one A4 page. Bring it to the audit committee as a discussion item. Ask: 'Which of these should IA always monitor?'

02 Next AC meeting: 0 extra hours

Ask the survivability question

Add one question to your verbal AC update: 'What risk, rated low probability, would damage us most if we were wrong?' Then stop. Let the silence work. Output: one answer from the committee you did not have before.

03 This week: 30 minutes

Run Level 1 risk sensing right now

Open Claude. Upload your risk register. Paste the prompt from slide 59. Read the output. Output: a list of gaps your register does not cover. Share it with one colleague. If you find one new risk, the process has already paid for itself.

“Thailand is not a seismic zone.

Until March 28, 2025, it was.”

What scenario are you not ready for?

OpenQ

Thank you

Noratip Dhanasarnsilp

tong.noratip@bitkub.com // noratip@qopen.co

CISA · CISSP · ITIL Expert · ISO 27001 / 27701 Lead Auditor

Clinic IA 5/2026 | IIAT × ISACA Bangkok Chapter | 27 June 2026

Research sources

- WEF Global Risks Report 2025
- Deloitte 2026 Hot Topics for Operational and IT IA (UK)
- EY 2026 Audit Committee Priorities Report
- KPMG 2026 Internal Audit Key Risk Areas
- PwC Risk Agenda 2026
- IIA Risk in Focus 2026 (4,073 respondents · 131 countries
- OECD Economic Surveys Thailand 2025
- ECB Financial Stability Review May 2026