

MITSUBISHI ELECTRIC · PRACTICAL CYBER × DACTA × APEXAGEN

The BMS Cyber-Sales Training

Day 2 — The AI Co-Pilot & The Pitch

Practical
Cyber

DACTA



APEXAGEN
TECHNOLOGY

AI does the intellectual labour. You keep the accountability labour.

Reading, sorting, drafting → the AI. The numbers, the commitments, the signed bid → you. The 98/2 split.

The AI Amplifier

- AI doesn't create new problems — it **amplifies** the ones you have
- A messy tender process + AI = a faster, more confident mess
- A disciplined one + AI = a force multiplier
- So we build the discipline first, then add the speed

98%

Deterministic — your CRM, figures, systems of record

2% · AI at the edge

drafting, summarising, sorting — never the record

AI does the intellectual labour; **you keep the accountability.**

Back to bms.ethanseow.com

- Check in for Day 2
- Today's drills: **Safe-AI, Prompt Builder, Tender Ledger, Pitch Builder**
- **Download** both days' slides to keep



MODULE 5

Which AI? Tools & Policy

Understand the machine well
enough to trust it appropriately
— then set the one rule.

How you'll use AI here — a web chat, by hand

Any chat works

ChatGPT, Claude or Gemini — the free web version is fine.

Manual, not auto

You paste in, it drafts, **you copy out and verify** — no auto-pilot.

The one caution

A free web chat may **keep what you paste** — so never paste a real client's live tender.

No agent needed

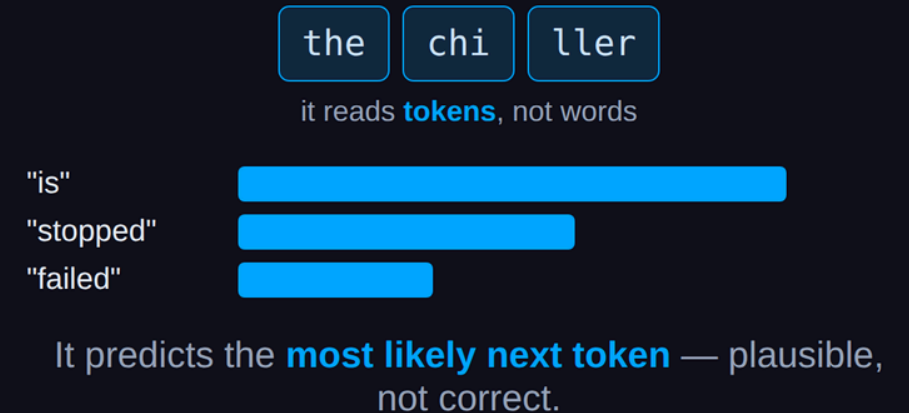
Everything runs in a plain chat box — no special "agent" tools.

You are the operator. The chat is a fast intern. You drive, and you sign.

How an LLM works (1/3) — it reads tokens, not words

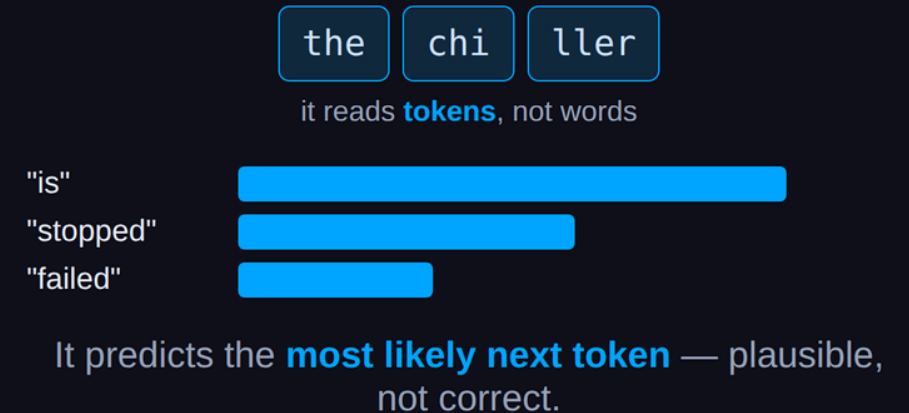
- Your text shatters into **tokens** — numeric fragments, not words or letters
- It never sees spelling — so it can't reliably **count, spell, or do exact maths**
- **Why** it hands you a confident, slightly-wrong clause or policy number

Not a database, not a calculator.



How an LLM works (2/3) — it predicts the next token

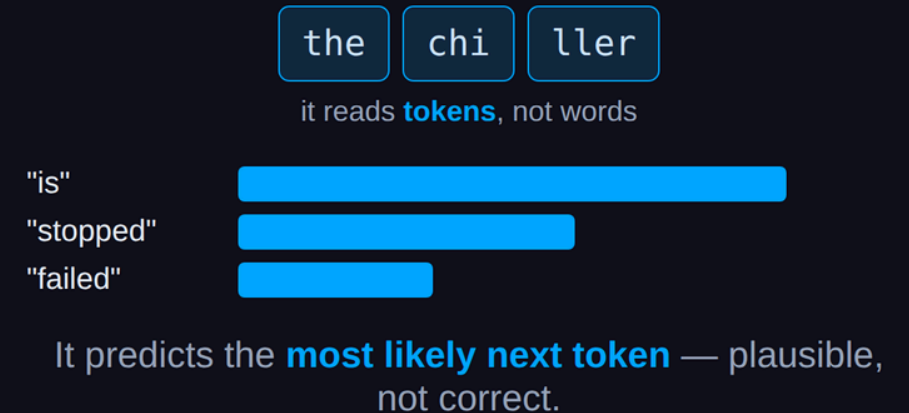
- Every reply is the **most likely next token**, over and over — weighted dice
- Like a **brilliant intern**: knows everything in general, nothing about **your** clients
- No memory, no live data — unless **you** give it both



How an LLM works (3/3) — so it hallucinates

- A **hallucination** is confident, wrong output — the machine has no ground truth
- It isn't lying; it produced the most **plausible** text, not the most **correct**
- Fluency hides the error

This is the whole reason for "verify every number" — it's how the machine works.



Risk — the Eloquence Trap

What it is

Polished AI wording that **sounds** authoritative but may be false.

Top tender risk

The #1 AI risk in a tender.

Say it

Say: "sounds right ≠ is right — check the claim, not the confidence."

AIR CANADA · 2024 · THE ELOQUENCE TRAP

source: [cbsnews.com](https://www.cbsnews.com) →

An AI-invented promise became binding

- The chatbot **invented a refund policy** that didn't exist
- A tribunal held the airline **to the AI's promise**
- An AI-fabricated compliance claim in your bid is **your** commitment

[MoneyWatch](#)

Air Canada chatbot costs airline discount it wrongly offered customer

By [Megan Cerullo](#)

February 19, 2024 / 1:05 PM EST / CBS News

[Add CBS News on Google](#)

Air Canada is being held responsible for a discount its chatbot mistakenly promised a customer, the Washington Post [reported](#).

The airline must refund a passenger, Jake Moffat, who two years ago purchased tickets to attend his grandmother's funeral, under the belief that if he paid full price, he could later file a claim under the airline's bereavement policy to receive a discount, according to a [ruling](#) by Canada's Civil Resolution Tribunal (CRT).

He didn't invent the idea, rather a support chatbot with which he communicated on Air Canada's website provided him the false information, ultimately costing the airline several hundred dollars. The tribunal's judgment could set a precedent for holding businesses accountable when relying on interactive technology tools, including generative artificial intelligence, to take on customer service roles.

In November 2022, Moffat spent over \$700 (CAD), including taxes and additional charges, on a next-day ticket from Vancouver to Toronto. He made the purchase after being told by a support chatbot on Air Canada's website that the airline would partially refund him for the ticket price under its bereavement policy, as long as he applied for the money back within 90 days, the tribunal document shows. Moffat also spent more than \$700 (CAD) on a return flight a few days

Risk — Shadow AI

The habit

Staff quietly using **unapproved** tools because they're faster.

The exposure

The danger: client data pasted into a consumer tool.

Say it

Say: "we use the sanctioned tool — never a free one for client data."

SAMSUNG · 2023 · THE IRREVERSIBLE LEAK

source: techcrunch.com →

Confidential data into a public tool

- Engineers pasted **confidential code** into ChatGPT to debug it
- It left the company's control — potentially into training data
- Your equivalent: a client's fact-find, drawings, or tender

THE IT NERD
Straight Talk About Information Technology From A Nerd
Who Speaks English

« Twitter's Purge Of Legacy Verification Checkmarks Has Made No Difference To The Number Of Twitter Blue Subscribers
Employer Disqualifies Any Candidate That Doesn't Have An iPhone.... WTF? »

Samsung Bans Internal Use of AI After ChatGPT Source Code Leak

On Monday, Bloomberg reported that Samsung notified staff of a new policy banning employee use of AI tools in response to discovering in April that its engineers had accidentally leaked internal source code by uploading it to ChatGPT.

The company is concerned that once data is transmitted to AI platforms it is then stored on external servers, difficult to retrieve and delete, and is then available to other users, according to the document disclosed to Bloomberg.

Furthermore, last month Samsung conducted an internal survey which revealed that 65% of respondents believe the use of AI tools poses a security risk.

Meanwhile, the company claims to be creating its own internal AI tools for translating and summarizing documents as well as for software development. It's also working on blocking staff's ability to upload proprietary information to external services.

Other companies that have either banned or restricted the use of AI tools include JPMorgan Chase & Co., Bank of America Corp., Citigroup Inc., Deutsche Bank, Goldman Sachs, and Wells Fargo.

Roy Akerman, Co-Founder & CEO, *Rezonate* had this comment:

"The wide adoption of AI language models is becoming widely accepted as a means of accelerating delivery of code creation and analysis. Yet, data leakage is most often a by-product of that speed, efficiency, and quality. Developers worldwide are anxious to use these technologies, yet guidance from engineering management has yet to be put in place on the do's and don'ts of AI usage to ensure data privacy is respected and maintained."

"The aspect of AI consuming all input as source material for others queries presents a black box of uncertainty as to exactly how and where a company's data would end up and completely upends the tight data security at the heart of most all companies today."

"Blanket restrictions are not a permanent solution and will only limit an organization's visibility to this problem. Instead, increased control, with education of developers on the cause and effect of using these tools for code reviews, code optimization, debugging and syntax will help harness the technology for the betterment of the organization."

Clearly there are advantages and benefits to using AI, but there are risks as well. Companies need to weigh those risks so that they aren't inadvertently creating a situation where AI does more harm than good.

Share this:



Like this:

☆ Like

Be the first to like this.

Related

Research: S&P 500's software companies face mounting AI-related threats
July 15, 2025
In "Commentary"

The Galaxy Tab S11 turns AI into a real productivity tool
January 27, 2026
In "Commentary"

Guest Post: Navigating the security and privacy challenges of large language models
November 7, 2023
In "Commentary"

This entry was posted on May 8, 2023 at 8:30 am and is filed under [Commentary](#) with tags [Samsung](#). You can follow any responses to this entry through the [RSS 2.0 feed](#). You can [leave a response](#), or [trackback](#) from your own site.

One Response to "Samsung Bans Internal Use

13

Practical Cyber × DACTA × Apexagen · BMS Cyber-Sales Training · Day 2

Privacy & Cookies

We and our advertising partners store and/or access information on your device and also process personal data, like unique identifiers, browsing activity, and other standard information sent by your device including your IP address. This information is collected over time and used for personalised advertising, advertising measurement, audience research and services development specific to our ads program. If this sounds good to you, select "I Agree" below. Otherwise, you can get more information, customize your consent preferences for 101 different ad vendors, or decline consent by selecting "Learn More". Note that your preferences apply only to this website. If you change your mind in the future you can withdraw your consent anytime by visiting the Privacy link displayed under each ad. One last thing, our partners may process some of your data based on legitimate interests instead of consent but you can object to that by choosing "Learn More" and then disabling the Legitimate Interests toggle under any listed Purpose or Partner.

Where your data goes — consumer vs enterprise

Consumer

May **retain and train** on your input — **irreversibly**.

Enterprise

With a no-training agreement, **isolates** what you paste.

Irreversible

You cannot un-train a model — a leaked spec is leaked forever.

The one rule

This is the mechanism behind the one rule.

Risk — the vendor opt-out trap

The default

An existing tool may quietly **turn AI training on by default.**

The effect

Your inputs get learned from unless you find the setting.

Check first

Check the setting before any client data flows.

Risk — human-in-the-loop

Never auto-act

AI must never **auto-send** to a client or **auto-commit** anything.

Human reads

A human reads every outbound message.

The gate

The sign-off gate is your accountability made structural.

What AI CAN do (advisory)

Summarise

A tender to a one-page action list.

Sort

Clauses into the Pillars.

Draft

Rationale, narrative, a pitch skeleton.

Surface

Red-herrings and ambiguities.

What AI CANNOT do without you

No claims

Submit a compliance claim or posture.

No commitments

Commit a price, date, or scope.

No sending

Send anything to a client.

No representations

Make a representation you'd be held to.

The right column is the 2% — and it's the job.

The one rule

Client data

Never put a client's confidential information into an **unsanctioned** tool.

Verify claims

Never submit an AI claim or number a human hasn't verified.

Everything else

Everything else is a corollary of this.

The Safe-AI checklist

Sanctioned tool

Enterprise, no-training — not a free one.

No client IP

Never in a consumer tool.

Verify everything

Every figure & compliance claim before it enters a bid.

You own it

No auto-send; you own every posture.

Safe / Sign-off / Never

- Eight real scenarios — make the call before AI touches a client tender
- Then write your team's 3 DOs and 3 DON'Ts

Web: [Safe-AI Decision](#) — merged into the class checklist you keep.

MODULE 6

Prompt Engineering — CAGE

Asking well is a sales skill now.
CAGE frames the job so the AI
is useful and safe.



CAGE — the anatomy of a good prompt

- C Context · A Align · G Goals · E Examples
- A weak prompt ("summarise this tender") is a gamble
- CAGE turns AI from a gamble into a tool you can defend

The skill is asking well.

C · Context

all of it — who, the client, the rules

A · Align

play back the plan, wait for your go

G · Goals

what "done" is + the guardrails

E · Examples

one good, one bad

Frame the job well — the skill is **asking well**.

C · Context

- **All of it** — who you are, the client, the situation, the rules that bind the task
- In your work you can't separate "their board just had a breach scare" from "their CCoP audit lands in March" — both shape the pitch
- Vague in, vague out

C · Context

all of it — who, the client, the rules

A · Align

play back the plan, wait for your go

G · Goals

what "done" is + the guardrails

E · Examples

one good, one bad

Frame the job well — the skill is **asking well**.

A · Align

- Make the AI **play back its plan and wait** for your go before it drafts
- The highest-leverage letter — 30 seconds saves 30 minutes
- Correct the approach **before** the output, not after

C · Context

all of it — who, the client, the rules

A · Align

play back the plan, wait for your go

G · Goals

what "done" is + the guardrails

E · Examples

one good, one bad

Frame the job well — the skill is **asking well**.

G · Goals

- What "done" looks like — the decision it serves + success criteria
- The **guardrails live here**: Must NOT invent, Must FLAG ambiguous
- Your compliance line becomes a **criterion**

C · Context

all of it — who, the client, the rules

A · Align

play back the plan, wait for your go

G · Goals

what "done" is + the guardrails

E · Examples

one good, one bad

Frame the job well — the skill is **asking well**.

E · Examples

- One or two samples of the output shape you want — one good, one bad
- The **cheapest accuracy gain** in the building
- The model copies the shape you show it

C · Context

all of it — who, the client, the rules

A · Align

play back the plan, wait for your go

G · Goals

what "done" is + the guardrails

E · Examples

one good, one bad

Frame the job well — the skill is **asking well**.

Why CAGE works — the mechanism

Context

More specific context narrows the model's output toward **your** answer.

Align

Inserts a checkpoint before it commits.

Goals & Examples

Goals turn constraints into testable criteria;
Examples anchor the shape.

Deliberate

Used deliberately, not as a ritual.

CAGE applied — weak vs CAGE

WEAK

- "summarise this tender"
- No context, goal, example or guardrail
- A vague wall of text; may invent clauses

CAGE

- C BMS bid, rail CII client, §42
- A play back your plan, wait
- G sort to Pillars; quote the clause; flag [UNVERIFIED]
- E one sample row

Same 60 seconds. An order of magnitude better — and defensible.

Tokens & the prompt economy

Context is a feature

Long context is a **feature** —
paste the **de-identified** tender
in, don't summarise it.

More is better

The model gets better with
more context, not worse.

Your scarcity

Your scarcity is your time, not
the AI's tokens.

ARCH — the loop you run by hand

- Action · Reasoning · Contextual-check · Horizon — what a **governed agent** would do on its own
- In a plain web chat there is no agent, so **you run the loop**: ask, read its reasoning, check it against the real §42, decide the next step
- CAGE frames the job; ARCH is the rhythm; **you are the loop**

A · Action

the agent acts



R · Reasoning

it shows its working



C · Contextual check

against the real tender



H · Horizon

hands the decision back to you

CAGE frames it · **ARCH** runs it · you decide. ↻

Build a CAGE prompt

- Assemble Context · Align · Goals · Examples — load a recipe, copy it out
- Run it on a §42 extract and **verify two claims by hand**

Web: [Prompt Builder \(CAGE\)](#) — speed is real; verification is non-negotiable.

MODULE 7

Tender Breakdown → The Posture Ledger

From "decode it" to the
disciplined, auditable way
professionals respond to it.



Tender Triage

First pass

First pass: sort every clause into
Must-have / Nice-to-have /
Risky red-herring.

Use the Pillars

Structure by the Pillars from
Day 1.

Mind the traps

The traps are where deals are
lost.

Triage — Must-have

What they are

The genuine, core requirements you meet with evidence.

Lead with these

Lead with these — they're your Comply column.

Examples

e.g. MFA on remote access; segmentation; asset inventory.

Triage — Nice-to-have

What they are

Real but non-critical — a dashboard, a mobile app.

The balance

Don't over-invest; don't ignore.

How to answer

Comply where cheap; note where not.

Triage — the Red-herrings (the traps)

"100% secure"

Indefensible; never warrant it.

"CCoP-certified product"

CCoP binds the operator, not a product.

Antivirus on every controller

Impossible on a DDC.

| An eager "Comply" here loses the bid.

The four disciplines of a winning bid

1 Anchors

Write each posture once.

2 Posture Ledger

Clause by clause.

3 Round-trip

Answer on the client's own file.

4 Phrasing discipline

Never overclaim.

Discipline 1 · Anchors

Write once

A standard position held across every bid, written **once**.

Posture library

A living **posture library** of true, reusable answers.

Change once

Change the anchor, not the 30 proposals that quote it.

Discipline 2 · The Posture Ledger

- Every clause: **Posture + Rationale + Evidence + Anchor**
- Evidence = a cert, a standard config, a past project
- It's what survives an audit

CLAUSE	POSTURE	RATIONALE
42.5 MFA	Comply	per-engineer MFA, jump host
42.9 AV on all	NC + alt	DDC can't run agents → segmentation
42.5.1c "eight(12)"	Clarify	a drafting slip — ask, don't assume

Every clause: **posture + rationale + evidence + anchor.**

Discipline 3 · Round-trip on the client's file

Their document

Answer on **their** document — same clause IDs, same order.

No reformat

Never hand back your own reformat — it reads as evasion.

Canonical container

Compliance returns in the prime's canonical container.

Discipline 4 · Phrasing discipline

The check

A **forbidden-phrasings** check before anything goes out.

Banned phrases

No "fully CCoP-certified", no "mandatory" to a non-CII client, no "100% secure".

Trace or pull

Every claim traces to evidence — or it's [UNVERIFIED] and pulled.

The four postures

- **Comply** — we meet it, with evidence
- **Not-comply (+ alternative)** — the honest, often-winning answer
- **Clarify** — ambiguous or mis-stated; question it
- **N/A** — genuinely doesn't apply

CLAUSE	POSTURE	RATIONALE
42.5 MFA	Comply	per-engineer MFA, jump host
42.9 AV on all	NC + alt	DDC can't run agents → segmentation
42.5.1c "eight(12)"	Clarify	a drafting slip — ask, don't assume

Every clause: **posture + rationale + evidence + anchor.**

Honest "Not-comply + alternative" beats overclaimed "Comply"

The reality

A DDC controller **can't run an AV agent** — "Comply" is a lie an auditor catches.

Strong answer

Strong: "**Not-comply on field controllers; segmentation + monitoring + whitelisting as compensating controls.**"

The disqualifier

Overclaiming is the disqualifier.

Honesty with a better alternative wins.

Worked row — §42.9 (antivirus on all devices)

Pillar

Devices · **Posture:** Comply — with **NC + alt** on field controllers.

Rationale

DDC/PLC can't run agents → compensating controls.

Anchor

ot-compensating-controls

Worked row — §42.5 (the "eight (12)" slip)

The slip

The clause literally prints "eight (12) characters" — a real drafting contradiction.

Posture

Clarify — don't blindly Comply to a contradiction.

The expert move

The expert move: you noticed the slip everyone else glossed over.

The real agentic tender-framework — the ceiling

The parts

Anchors in **version control** · a **self-linting ledger** · **lodgement onto the client's exact file.**

The linter

The linter **refuses** to lodge an overclaim or an open question.

Who does what

AI reads, proposes, drafts; the framework enforces; a human decides.

| The ledger you built by hand is its portable version.

Build the Posture Ledger on the real Section 42

- Triage → sort to Pillars → set a posture per clause, on the real tender
- Find the honest **NC + alternative** and the **Clarify**

Web: [Tender Response Ledger](#) — the day's biggest hands-on.

MODULE 8

The Resilience Roadmap Pitch

Turn the breakdown into a business case, and defend it to a board.



The five pitch blocks

- **1 Frame (SRP) → 2 Crown Jewels → 3 The 80/20 → 4 Compliance position → 5 Value Protected**
- Reverses the vendor default: lead with what the owner protects, not features
- Close on the business outcome

1 · Frame (SRP)

what the owner is protecting

2 · Crown Jewels

what matters in this building

3 · The 80/20

the vital few controls

4 · Compliance position

honest, CCoP-precise

5 · Value Protected

the business outcome, with numbers

Lead with what they protect — **close on value.**

Block 1 · Frame (SRP)

The opening

Open on what the owner is protecting: **safety, continuity, value.**

Not features

Not a feature list — the stakes they already care about.

Safety first

Safety first, always.

Block 2 · Crown Jewels

Be specific

Name what actually matters in **this** building.

The assets

Fire, the head-end, the controllers, the remote-access layer.

Process not paper

Protect the process, not the paperwork.

Block 3 · The 80/20

Vital few

Lead with the **vital few** controls, not a wish-list.

The controls

Inventory · segmentation · vendor access · isolation · degraded-mode.

Depth that matters

Depth where it moves the risk.

Block 4 · Compliance position

Honest & precise

Honest and CCoP-precise —
no overclaim.

The line

"We operate to CCoP-grade
controls; here's exactly what
binds you"

Precision

Precision is the credibility.

Block 5 - Value Protected

The close

Close on the **business outcome** — with numbers.

The line

"We protect S\$X/hour of cooling-dependent operations"

Number not adjective

The number, not the adjective.

Quantifying Value Protected

- **Downtime cost/hour** — a data hall, a hospital ward
- **Safety/liability** exposure · **insurance** premium & eligibility
- **Asset value / green-mark** rating

Value Protected is a number. Find the owner's number and protect it.

Downtime cost / hour \$\$\$

Safety & liability exposure \$\$\$

Insurance premium & eligibility \$\$

Asset value / green-mark \$\$

= Value Protected (a number)

Find the owner's **number** and protect it.

Value — downtime cost

The question

What does an hour of lost cooling / lifts / access cost this owner?

The scale

Data halls and hospitals:
measured in millions/hour.

Anchor to it

Anchor the pitch to that figure.

Value — safety & liability

The cost

The legal and human cost of a life-safety failure.

The exposure

Insurance, regulatory, reputational exposure.

The reduction

The controls reduce a quantifiable liability.

Value — insurance & asset value

Insurers require

Insurers increasingly require cyber controls to cover a building.

Rates & rents

A secure, well-run building **rates and rents** better.

Asset value

Security protects the asset's valuation.

Decision Survivability — defend every claim

The board's job

The board's job is to **pressure-test** you — especially for overclaims.

Trace every claim

Every claim traces to a reason or evidence you can defend.

Who wins

The winners aren't the ones who "Comply" to everything.

Could you defend this after something went wrong? Then you've earned the room.

Handling the board probe

Expect it

Expect the probe on your **boldest claim**.

Don't bluff

Don't bluff: "**Not-comply as written; here's our stronger, defensible alternative.**"

Bonus

Bonus for **clarifying** a mis-stated requirement.

Build & deliver the pitch

- Build the Roadmap from your \$42 ledger — five blocks, Value Protected with numbers
- Pitch to the mock **Board of Building Owners**; expect the probe

Web: [Resilience Roadmap Builder](#) — scored on honest postures, CCoP precision, Decision Survivability.

You can read a cyber tender, translate it, respond honestly, and defend every word.

That's the Practical Cyber BMS Sales Consultant — the salesperson your clients can't fool, and can't ignore.