The Tender-Prep Kit

What to gather *before* you point AI at a tender

A BMS Cyber-Sales Training leave-behind · Practical Cyber × DACTA × Apexagen

---

Your AI answer is only as good as the material you feed it. The AI is the *super-powered intern* — brilliant in general, but it knows **nothing about your products or your house positions** until you tell it. So on its own it can *triage* a tender and *spot the traps*, but it can't say *"we comply on remote access — our controllers do native MFA"* unless you've handed it the facts.

> **The tender is the question. This kit is the answer key.** Gather it once, reuse it on every bid. Miss it, and you get a fast, confident, *thin* answer.

This kit is your **tendering Vault** — the reference base you load into an AI Project (see the *Projects & Anonymise Guide*) or feed to a private agent.

The five buckets

1 · Product & capability data *(the one people forget)*

Why: the AI can't answer **Comply / Not-comply** without knowing what the kit actually does.
- ☐ Controller / HMI / head-end **datasheets**
- ☐ A **cyber-capability matrix** per product line: MFA? logging/audit? segmentation-ready? role-based access? (yes / no / with-add-on)
- ☐ **Firmware support & end-of-life** lifecycle list (for the "unsupported firmware" clauses)
- ☐ **Compensating-controls** list — what you offer where a control is *impossible* (DDC can't run AV → segmentation + monitoring + whitelisting)
- ☐ Your **standard / reference architecture** diagram (how you segment a BMS)

2 · Standard postures — your Anchors

Why: the highest-leverage reuse is the *true, defensible answer written once* — change the anchor, not the 30 proposals that quote it.
- ☐ **Posture library**: the reusable answer to each common clause — remote access, segmentation, asset inventory, backups, hardening, incident support
- ☐ **Forbidden-phrasings** list — never "100% secure", never "CCoP-certified *product*", never "mandatory" to a non-CII client
- ☐ A couple of **past winning responses** to show the shape of a good answer

3 · Standards & mappings

Why: so the AI can map *clause → standard → pillar → your posture* instead of guessing.
- ☐ The **CCoP clause map** (the Field Advisory)
- ☐ The **7 Pillars** one-pager
- ☐ **IEC 62443** zones-and-conduits reference

4 · Client & situational intel

Why: Context is the human + institutional picture, not just the clauses.
- ☐ Is the client **governed (CII)**? which sector?
- ☐ The **live situation** — board just had a breach scare, CCoP audit due, incident history
- ☐ **Past bids** to this client or similar buildings (what won, what lost)

5 · Scope & guardrails

Why: so it doesn't quietly over-commit you.
- ☐ Where **your scope ends** — the BMS segment vs. the SI's / owner's infrastructure
- ☐ Rough **pricing / effort** guardrails
- ☐ **Sign-off rule** — who approves a posture before it goes in a bid

---

The bottom line

Tender + buckets 1–3 = a comprehensive, defensible draft. Buckets 4–5 make it *this client's* answer and keep you honest. Without 1–3 the AI free-wheels — and it will "Comply" its way straight through the traps with perfect grammar (the **Eloquence Trap**).

Who owns the Vault? Start with your own mini-kit. The real win is when **ME builds one shared tendering Vault** — every bid pulls the same true Anchors, and you *change the anchor once* instead of re-arguing it in 30 proposals. That's the Institutional Vault, applied to sales.

> You still own the 2%: the AI compares and drafts; **you set the posture and sign it.** *(Decision Survivability — could you defend this after something went wrong?)*

---

*Pairs with the **Projects & Anonymise Guide**. Practise on the fictional tender at **bms.ethanseow.com/practice-tender.html**. Practical Cyber × DACTA × Apexagen.*