Executive Education — The Senior Professional's Standing Course

You have spent a decade or more winning work. The executive tier exists so that the next decade is not just more of the same — so that your judgment, your portfolio, and the organization you lead keep getting sharper while the rulebook, the market, and the machine corpus keep changing around you.

Level: executive · **Audience:** senior GovCon professionals (typically 10+ years), pursuit and BD executives, firm owners, C-suite · **Companions:** course/forever-learning-program.md, literacy/industry-standard-canon.md.

This module is the standing curriculum for the working professional who has already mastered the craft and now needs the **strategic altitude**. It is not a re-run of the master's tier. It is the tier where your own experience becomes the teaching asset, where the rulebook's changes are tracked as they happen, and where the discipline's hardest questions are discussed among peers who have all felt them.

1. Who this is for — and what it is not

This is for the **pursuit leader**: the capture director who has run fifty captures, the BD vice president accountable for a portfolio, the firm owner deciding which markets to enter, the contracts executive who has survived a DCAA audit, the CEO whose win rate is the company's growth engine.

Three things this tier is *not*:

- **Not a fundamentals course.** If a learner cannot already build an ORBITAL, score a pursuit, and read a solicitation at professional altitude, they belong in the graduate tier, not here.
- **Not a product training.** This tier is stack-agnostic with the rest of the curriculum (doctrine/08). No session requires operating any system of record, any AI layer, or any tool — the executive's own firm is the sandbox, and the doctrine is the lens.
- **Not a spectator series.** Every format in this module puts the executive to work: analyzing a real NOFO, defending a gate charter, peer-reviewing a live capture. The value is in the doing.

The entry test is honest and short: *can you take a real, published-and-closed competition and produce a win strategy, a pricing posture, and a gate plan that would survive a skeptical board?* If yes, this tier is yours.

2. Executive master classes

The **executive master class** is a short, high-intensity encounter — a half day to two days — on one strategic topic at full altitude. Each class is concept-first: it teaches the *idea* at executive depth, uses a real, published-and-closed case as the working object, and sends the executive home with a decision artifact. The recurring catalogue:

Master class	The strategic question	The decision artifact
The Strategy of the Portfolio	How do you allocate scarce capture capacity across a pipeline of opportunities — and when do you say no?	A portfolio allocation and a no-bid discipline for the next quarter.
The Economics of the Win Rate	What does a 5-point improvement in win rate do to revenue, and what is it worth paying to achieve?	A win-rate economics model for the learner's own firm.
Gate Governance at Scale	How do you design the gates, decision rights, and ethical walls for an organization that pursues dozens of opportunities at once?	A gate charter for the learner's organization.
The Shape of the Market	Where is the money moving — and which markets, vehicles, and set-asides will reward a firm five years from now?	A market-position thesis, argued against a peer panel.
Price-to-Win as Strategy	How does a firm *choose* its pricing posture rather than react to the competitor band?	A pricing-posture doctrine for the learner's firm.
Team, JV, or Solo	When does the business structure of a pursuit — teaming, joint venture, mentor-protégé — become a strategic advantage rather than a compliance exercise?	A teaming-posture framework for the next pursuit.

A master class is graded only by the quality of the artifact and the argument. The class's real assessment is whether the executive's firm acts on it.

Worked session plan — The Economics of the Win Rate

The strategic question the class answers: **what does a 5-point improvement in win rate do to revenue, and what is it worth paying to achieve?**

The class is a half day (three and a half hours, with a break), run for a working group of 6–12 executives. It is concept-first: the arithmetic is taught as a durable model the executive can rebuild on any spreadsheet or on paper, never as a product workflow (doctrine/08).

Learning objectives. By the end of the class, an executive can:

- Build a win-rate economics model for their own firm: revenue per win, pursuit cost per bid, conversion assumptions, and the sensitivity of revenue to a one-point change in win rate.
- Distinguish the three levers the model exposes — win more, bid smarter (raise conversion without more bids), and spend more efficiently (the marginal cost of the next pursuit) — and say which is cheapest for their firm.
- Price the investment question: what is it worth paying for a pipeline, a capability, or a price posture that moves the win rate by a defensible amount — and when is the answer "less than you think, because the base is small"?

Agenda (half day).

Time	Block	What happens
0:00–0:20	Open the question	The facilitator poses the strategic question and runs the room's first guesses: "What is your firm's win rate, and what would a 5-point improvement be worth?" Executives log their priors before any arithmetic.
0:20–1:00	Teach the model	The win-rate economics model at concept level: revenue per win, the pursuit-to-award conversion funnel, the cost per bid, and the identity that ties them — revenue = (opportunities × conversion × value per win). The three levers and the arithmetic of each.

Time	Block	What happens
1:00–1:50	Build the firm’s model	Each executive builds the model for their own firm with real (or, where numbers are private, clearly-labeled placeholder) inputs: current conversion, value per win, cost per bid. The class works individually, then in pairs.
1:50–2:05	Break	
2:05–2:50	Sensitivity and the investment question	The room reads its models at altitude: which lever moves revenue most for each firm, and what the marginal cost of moving it would be. The facilitator introduces the framing — *value of a point of win rate* as a durable number the executive can carry into any budget conversation.
2:50–3:20	Defend a recommendation	Each executive states one sentence: “the highest-value lever for my firm is X, and it is worth paying up to Y for it.” Peers attack the reasoning the way a board would.
3:20–3:30	Close	The decision artifact is committed: the win-rate economics model, the one-lever recommendation, and the number the executive will take to their next budget decision.

Materials. The doctrine files doctrine/03 (the pipeline — conversion as a funnel) and doctrine/05 (the score and threshold) at concept level; the class's one-page win-rate-economics model sheet (the identity, the three levers, the sensitivity template); each executive's own firm numbers (or clearly-labeled placeholders — the class follows the curriculum's honesty standard: no invented specifics passed off as real).

Discussion prompts.

- "Which of the three levers is cheapest for your firm right now — and why do firms default to the most expensive one?"
- "When is a 5-point improvement in win rate worth *less* than the room's intuition says — what makes the base small?"
- "Defend a position: is the marginal pursuit worth more as a win (revenue) or as a learning event (past performance)?"
- "What would the model have to show for you to *stop* investing in win-rate improvement — and is that a number your firm has ever actually computed?"

Facilitator note. The class's failure mode is arithmetic theater — executives nodding at a model they will not rebuild. Guard against it two ways. First, require the *sensitivity* step: no executive leaves without naming which lever moves their revenue most, because that is the number they will actually use. Second, keep the safe-exercise rule: real firm numbers are welcome and the room is confidential, but any executive who cannot share real numbers works with clearly-labeled placeholders — nothing invented is presented as fact. The facilitator's hardest job is the defense block: let peers push until a recommendation either survives or gets revised in the room.

3. The pursuit-leadership studio

The **pursuit-leadership studio** is the flagship format of this tier — a recurring working session where executives bring their *own* real work into the room and put it under peer discipline. It is the executive version of the color-team cadence (literacy/how-color-teams-work.md), applied to the leadership layer rather than the document layer.

Each studio cycle follows the pipeline's own discipline (doctrine/03, doctrine/09):

1. **A live capture, peer-pressed.** An executive presents a pursuit currently in their pipeline — the win strategy, the discriminators, the teaming, the pricing posture. Peers attack it with the same independence a red team applies to a proposal: Is the position real? What would the incumbent say? Where is the ghost theme hollow?
2. **A gate charter, designed and defended.** An executive brings an organizational-governance problem — a gate that keeps getting overridden, a bid/no-bid call that has no owner, an ethical wall that is leaking. The

studio designs the fix and defends it.

3. **A debrief, mined.** The studio reads a real loss — a debrief from a published-and-closed competition — and asks what the doctrine should have taught the team that did not learn it. The lessons feed the curriculum's debrief-retrofit loop (refresh-and-recertification.md).

The studio's rules are the doctrine's own: **real documents only, safe exercises** (nothing live, nothing source-selection-sensitive), **specific findings, forward-only cadence**, and **independence** — a peer panel that tells you what you need to hear, not what you want to hear. The studio is where "mastery is asymptotic" becomes practical: there is always a peer who has seen a harder capture.

A studio session, minute by minute

The studio is a recurring working session — typically three hours, meeting monthly, with a fixed roster and a rotating presenter. The spine is the pipeline's own discipline: a real piece of the presenter's work is put on the table, peers press it with independent, specific findings, and the room closes on the changes the presenter will make.

Learning objectives. By the end of a studio session, a participant can:

- Present one piece of live work — a capture, a gate charter, a debrief — in the discipline's own structure, and defend it against peer pressure.
- Give findings the way a color team demands: specific (naming the page, the clause, the position), independent (not captured by the presenter's framing), and forward-only (improving the next version, not relitigating the last).
- Leave with a written set of commitments the presenter has accepted, owned, and dated.

Agenda (three hours, one presenter, one live object).

Time	Block	What happens
0:00–0:10	The ground rules	The facilitator re-states the studio contract: real documents only, nothing live or source-selection-sensitive, specific findings, forward-only, independence. The presenter's object is named.

Time	Block	What happens
0:10–0:40	The presenter's read	The presenter lays out the live object in the pipeline's structure: for a capture — the win strategy, the discriminators, the teaming, the pricing posture; for a gate charter — the gate ladder, the criteria, the owners; for a debrief — what was predicted, what happened, what was learned. No interruptions.
0:40–1:10	Silent findings	Peers write their findings individually before any discussion — the discipline's way of keeping findings independent. Each finding names the place and the question.
1:10–2:00	The press	Findings are read out, one per peer, round-robin. The presenter may clarify but not defend at length; the room's job is to surface the flaw the presenter cannot see. The facilitator keeps findings specific and forward-only — a vague critique is sent back for a page number.
2:00–2:30	The presenter's response	The presenter accepts, revises, or rebuts each finding aloud, and commits to the changes they will make.

Time	Block	What happens
2:30–2:50	The commitments	The room agrees the written commitment list: what changes, by when, owned by whom. The list is the session's trace (it feeds the cross-cohort case library and the debrief-retrofit pipeline).
2:50–3:00	Close	One-line takeaways from the presenter and two from the room; the next presenter and object are confirmed.

Materials. The presenter's real document (a capture plan, a gate charter, a debrief from a published-and-closed competition — never anything live or source-selection-sensitive); the studio's finding form (place, finding, suggested change, severity); the doctrine files doctrine/03 (the pipeline — the present structure) and doctrine/04 (gates — the governance frame) as the shared vocabulary.

Discussion prompts.

- "What would the incumbent say about this position — and have you priced that response into your strategy?"
- "Where is this gate's criteria written, and who owns it — or is this gate a meeting, not a gate?"
- "If this capture lost today, what would the debrief say the doctrine should have taught you — and does the current strategy teach it?"
- "What is the one assumption you are most attached to, and what evidence would change your mind?"

Facilitator note. The studio lives or dies on independence. Three discipline moves protect it: **silent findings first** (peers commit their findings before the presenter's charisma shapes them), **specificity enforced** (a finding without a page or a position is returned, not recorded), and **forward-only enforced** (the room never relitigates; it improves the next version). The presenter's response block matters as much as the press: the session is not over when the findings are read, it is over when the presenter has committed to changes. The facilitator also guards the safe-exercise boundary as the room's standing rule — nothing live, nothing source-selection-sensitive ever enters the studio. If a presenter drifts toward a live competition, the facilitator stops the session and redirects to the published-and-closed object.

4. Current-event roundtables: real NOFOs in the safe sandbox

The **current-event roundtable** is the discipline's newsroom. It meets on a cadence — monthly or quarterly — to put a *real, published* solicitation on the table and read it at full altitude.

The roundtable's discipline is worth stating plainly because it is what makes the exercise safe and useful:

- **Published and closed only.** The teaching object is always a solicitation that is either already closed or already awarded — from SAM.gov, Grants.gov, or a published agency forecast. The roundtable never analyzes a live competition, and no participant may bring source-selection-sensitive information into the room.
- **Read it as an evaluator.** The group reads the Section L / Section M structure the way a source-selection team would: what is being scored, what the discriminators would have to be, and which offerors the structure favors.
- **Read it as a strategist.** Then the group builds the win strategy it would have used — themes, discriminators, teaming, pricing posture — and defends it.
- **Read it as a teacher.** Finally, the roundtable asks what this NOFO teaches that the doctrine does not yet say. That answer is the seed of a debrief-retrofit (refresh-and-recertification.md).

The roundtable keeps the executive's reading sharp against a *moving* market. A professional who reads fifty real solicitations a year never loses the muscle of seeing a requirement the way the government wrote it.

5. Regulatory-change watch: FAR / DFARS / DCAA / CAS / TINA

The rulebook is a living document, and its changes are often the most important market signal an executive can track. The **regulatory-change watch** is the standing structure that tracks it:

Layer	What changes	Why the executive must watch it
FAR	Parts amended via the federal acquisition rulemaking cycle; new clauses; updated thresholds.	Changes the compliance spine of every pursuit — and sometimes the competitive field itself.

Layer	What changes	Why the executive must watch it
DFARS	DoD-specific supplements, including cybersecurity and supply-chain clauses.	DoD's share of the market moves with its rule changes.
DCAA	Audit guidance, cost-accounting expectations, and enforcement posture.	A change in audit posture can change how a firm prices and books costs.
CAS	Cost-accounting standards coverage and administration.	CAS coverage changes the cost structure of every negotiated price.
TINA	Truthful-cost-or-pricing-data requirements and thresholds.	Threshold changes alter which bids demand certified cost or pricing data.

The watch is *not* a legal briefing — the curriculum teaches concepts, and the executive's counsel owns the legal advice. The watch is a **strategic lens**: it reads each regulatory change the way capture doctrine reads a solicitation, asking what it does to the competitive field, the pricing posture, and the firm's position.

6. Peer review of live captures

Where the pursuit-leadership studio works on a *presented* capture, this thread is the **reciprocal service**: the executive serves as a peer reviewer for another organization's or another cohort's live work. It is the discipline's mentorship duty made concrete — the experienced professional lending their judgment.

Peer review follows the color-team contract: **independence** (the reviewer is not on the pursuit team), **specificity** (findings name the page, the clause, the position), and **forward-only** (the review improves the next version, it does not relitigate the last). It is also a **learning device for the reviewer**: nothing sharpens judgment like being asked to find the flaw in someone else's win strategy. Executives who review others' captures regularly report seeing their own next pursuit more clearly.

7. The executive's reading list

The discipline has a canon (literacy/industry-standard-canon.md), and the executive should *lead* from it rather than skim it. The standing reading list, kept current by the refresh cadence:

1. **Shipley's capture and proposal method** — the profession's operating vocabulary, in its latest edition. The curriculum paraphrases it in its own terms; the executive should be able to read the source and map it back.
2. **The APMP Body of Knowledge** — the field's shared framework. The senior professional knows the BOK well enough to *teach from it*.
3. **The FAR, in living form** — not memorized, but navigated: Part 15 (negotiated acquisition), the cost principles (Part 31), and the clause matrix. The executive reads the FAR the way a lawyer reads a code — to find the structure, not to recite it.
4. **Decision science and game theory at altitude** — the master's courses name the sources (course/graduate-competitive-strategy-game-theory.md); the executive deepens them: Kahneman and Tversky on judgment, Dixit and Nalebuff on strategy, the auction-theory canon.
5. **The debrief and protest record** — the most under-read public data in the field. Reading debriefs and protests teaches more about how source selection actually works than any textbook.
6. **The machine corpus** — the DNA corpus, read as doctrine. An executive who can read the rules and formulas the machine operates understands exactly what the "AI layer" is doing — and where to trust it.

The reading list is not a fixed syllabus; it is a **living canon** updated by the refresh-and-recertification cadence as Shipley publishes, the APMP revises, and the FAR changes.

8. How the executive tier connects to the ladder

The executive tier sits at the top of the practitioner arc and feeds the research tier in both directions:

- **Up into the fellowship.** The executive's live captures and debriefs are the raw material of the outcome-data record the post-doctoral fellowship studies. An executive who logs honest predictions and outcomes is doing the field's most valuable research service without a publication.
- **Down into the tiers below.** The executive is the alumni community's senior mentor (alumni-community.md), the master class is the graduate tier's guest-lecture source, and the current-event roundtable keeps the whole curriculum honest against a moving market.
- **Across the fellowship.** The tier maps to **The Ascension** (the Launchpad's executives' program) and, at its most senior, to **The Transcendence** (stewardship). The degree fellowship crosswalk (align/concept-to-system-map.md) keeps the mapping current.

Self-check

1. What is the difference between the executive master class and a graduate course — and why does the executive's *own firm* become the sandbox?

2. Why does the current-event roundtable use only *published and closed* solicitations, and what does that rule protect?
3. How is serving as a peer reviewer also a learning device for the reviewer?