Practice Bank — Answer Keys

Model answers for every self-check in the doctrine, every comprehension check in the modules, and every exercise in the practice bank. Answers are models, not scripts — a defensible answer that says the same thing in your own words is a good answer.

> **How to use this file.** Part A answers the three-question self-checks at the end of each doctrine chapter. Part B answers the comprehension checks at the end of every module session. Part C answers the numbered exercises in exercises.md. The **decide** and **build** exercises (tagged and) have *defensible* answers rather than a single right one — the model answer shows the shape of a strong defense.

Part A — Doctrine self-checks

Doctrine 01 — Where the Money Lives

1. What are the three main kinds of federal money, and what is the difference between a contract and a grant?

The three kinds are **contracts**, **grants**, and **other transactions / cooperative agreements**. A **contract** is the government buying goods or services *from you* — governed by acquisition rules (the FAR). A **grant** is the government funding *a mission you carry out* — governed by grants policy. The one-sentence difference: *a contract buys you; a grant funds your mission.* (Other transactions and cooperative agreements sit in between — faster, more flexible shapes for prototypes and partnered research.)

2. What four identifiers does a business need to be able to receive federal money, and where do you obtain them?

1. **UEI** — the 12-character Unique Entity ID; free from SAM.gov.
2. **CAGE code** — the 5-character Commercial and Government Entity code; assigned through the SAM.gov registration.
3. **NAICS code** — the 6-digit industry classification; you *choose* the codes that describe your business (multiple are allowed).
4. **SAM.gov registration** — the master eligibility record; you register at SAM.gov.

The memorization sentence: *register once in SAM.gov, get a UEI, get a CAGE code, and choose the NAICS codes that describe your business.*

3. Why do small-business programs like HUBZone and SBIR exist, and how do they change the competitive landscape for a small company?

They exist because Congress **required** a portion of federal dollars to reach small businesses — a deliberate policy decision that the federal market should not be a private club. They change the landscape structurally: a certified firm competes in a **smaller, friendlier pool** (set-asides restricted to that class), and programs like SBIR/STTR give small firms a *structured path* to research funding. For a small company, the right certifications are not paperwork — they are a competitive advantage.

Doctrine 02 — RFPs, NOFOs, and Solicitations

1. What are the six questions every solicitation answers, and which one is the most important to a winning proposal?

1. **What** is being bought or funded? 2. **Who** is eligible? 3. **How** will we be judged? 4. **When** are the milestones? 5. **How** must the response be structured? 6. **Where** must it go, and how?

The **evaluation criteria** (question 3) is the most important to a winning proposal — *a proposal is a document written to the evaluation criteria*, and everything else is decoration.

2. What is the difference between a compliance reading and a strategy reading of an RFP?

The **compliance reading** answers *"what is required?"* — mechanical and unforgiving; it produces the compliance matrix of every "shall," page limit, and format rule. The **strategy reading** answers *"what does the evaluator actually care about?"* — the agency's real intent behind the letter of the rules. Beginners read only for compliance and produce compliant, mediocre proposals; professionals read for both, in sequence.

3. Why is the "two-deadline" trap a risk, and what invisible third deadline does it hide?

Every solicitation has at least two deadlines — a **questions deadline** (after which questions get answered by amendment) and a **submission deadline** (after which the portal closes and late is rejected). The trap is treating them as one. The invisible third deadline is **your internal deadline** — days before the real one — that leaves room for review, fixes, and the mechanical act of submitting. A pursuit is not ready because the calendar says so; it is ready because the review gates say so.

Doctrine 03 — The Pursuit Pipeline

1. Name the pipeline stages in order, and say in one sentence what each one does.

1. **Sense** — watch the river of opportunities and recognize what is relevant.

2. **Qualify** — fast, honest filter: are we eligible, is it real, does it fit? (earliest bid/no-bid).
3. **Score** — estimate probability of win × value and compare to a threshold.
4. **Pursue** — turn the opportunity into a structured ORBITAL with an owner, money, and a clock.
5. **Design** — produce the proposal: technical, management, past performance, pricing.
6. **Audit** — fresh-eyes review (color team) before anything ships.
7. **Submit** — the compliant, reviewed package goes to the portal on time.
8. **Deliver** — do the work; delivery is the next proposal's past performance.
9. **Learn** — study the outcome and bake the lesson into the next cycle; repeat.

2. What is a gate, and why does a gate need its criteria written *before* the work begins?

A gate is a deliberate checkpoint where work pauses for a judgment call — with a trigger, criteria, a decision, and an owner. The criteria must be written **before** the work begins because that is what makes the gate governance rather than a rubber stamp: if the criteria are invented at the meeting, the gate is theater. Criteria written in advance make the decision passable or not — and keep the organization honest when adrenaline is high.

3. How are the GROWTH lifecycle, the DREAM approach, and the SuperNova strategy different accents on the same pipeline?

They are the same shape told from different points of view. **GROWTH** (Gauge, Research, Organize, Write, Target, Health) is the production arc of a *single pursuit*. **DREAM** (Discover, Relate, Empower, Advise, Motivate) is how you *engage the customer and the mission*. **SuperNova** is the *full arc* from a seed idea to an operating entity (DREAM → ORBITAL Plan → STARTUP → Ops → LAUNCH → GROWTH → Mission Control). Different accents, one pipeline — a student should recognize the pipeline in all three.

Doctrine 04 — Gates and Governance

1. What are the four parts of a gate, and which one keeps a gate from being a rubber stamp?

A gate has a **trigger** (work pauses by rule), **criteria** (written in advance, passable or not), a **decision** (proceed, proceed with conditions, or stop), and an **owner** (a person accountable). The **criteria written in advance** is the part that keeps a gate from being a rubber stamp — if the criteria are invented at the meeting, the gate is not governance.

2. Name the four gate words used in this course and what each one protects.

- **Bid / no-bid** — protects attention and strategy: *should this be a pursuit at all?*
- **Threshold** — protects money: *does this clear the score bar?*
- **Review gates** (color team) — protects quality: *does the proposal meet the criteria?*
- **Readiness gate** — protects reputation: *is this package ready to represent us?* (the last gate that can stop a bad submission).

3. How are private-side gates and public-side governance (law and policy) two halves of one system?

The private-side gates (our discipline: bid/no-bid, threshold, readiness) decide *whether and how we pursue.* The public-side governance (the FAR, grants policy, deadlines, evaluation criteria, certifications) decides *how the money is spent fairly.* Both keep the same transaction honest: a pursuer who respects the rules is a professional; a pursuer who games them is not merely unethical but legally at risk — and harms the trust the whole market depends on. One transaction, two halves.

Doctrine 05 — Scoring and Price-to-Win

1. What two numbers make up a pursuit score, and why does multiplying them matter?

Probability of win (pWin) and **estimated value**. Multiplying them produces a *comparable number*: a $200,000 pursuit at 60% pWin (score $120,000) beats a $2,000,000 pursuit at 5% (score $100,000) — and the score says so without an argument. Multiplication matters because it weights a pursuit by *both* likelihood and worth, so hope cannot hide a low probability or a trivial value.

2. What does a threshold gate do, and why is the *discipline* of the threshold more important than its exact value?

A threshold gate draws a line: above it the pursuit advances; below it the pursuit is dropped or parked. The **discipline** matters more than the exact number because the threshold forces the hard conversation *before* resources are spent and gives an organization the courage to say no to an exciting, un-winnable pursuit. If your threshold is honest, you will occasionally decline opportunities others chase — and you will be fine.

3. What is price-to-win, and why is a bid that is too low sometimes worse than a bid that is too high?

Price-to-win is the price that is **low enough to be competitive and high enough to be real** — set at the intersection of your cost floor, the competitive range, and what the evaluation rewards. A bid that is **too high** loses on price. A bid that is **too low** *wins a contract that bleeds you* — the worst outcome in the discipline, because you get the work and lose money doing it. A too-low bid is a business risk, not just lost profit.

Doctrine 06 — The ORBITAL Business Structure

1. What are the seven axes of the ORBITAL, and what practical question does each one answer?

Axis	Question
Objective	What are we trying to do, and what counts as success?
Resources	Who is on the team, and what do we bring?
Budget	What will it cost, and what is it worth?
Indicators	How will we know we are on track?
Transport	How do the pieces reach the customer?
Activities	What work is going to happen, and in what order?
Logistics	What has to be true for the work to happen?

2. Why is a structured pursuit more survivable than an unstructured one — what does the skeleton protect against?

An unstructured pursuit has a recognizable failure mode: the details live in someone's head, the budget in someone else's spreadsheet, the deadline in an email. When the capture lead gets busy — or leaves — the pursuit **disintegrates**. The ORBITAL protects against that by putting every important dimension in one place, owned and visible. It also anchors the proposal (objective), makes the score honest (budget), makes gates evidence-based (indicators), and makes the work schedulable (activities, logistics).

3. How does the ORBITAL connect to the wider Dream → ORBITAL → World lifecycle?

The ORBITAL is the *middle state* of the lifecycle: a Dream (seed idea) is structured into an ORBITAL (a pursuit with a skeleton, budget, owner, and deadline), and an ORBITAL that wins and delivers becomes a World (a fully realized enterprise). The ORBITAL is the mechanism by which ideas gain momentum — *structure first, document second.*

Doctrine 07 — The Lifecycle: Dream → ORBITAL → World

1. What are the three states of the lifecycle, and what moves a thing from one state to the next?

Dream (a seed idea: loose, cheap, no budget, no owner) → **ORBITAL** (structured: seven axes, budget, owner, deadline, momentum) → **World** (fully realized: revenue, track record, reputation, people). **Structure** moves a Dream to an ORBITAL (the seven axes); **momentum — winning and delivering** — moves an ORBITAL to a World. The direction of growth is always the same: structure creates momentum, and momentum creates mass.

2. How are the pipeline, the GROWTH stages, and the lifecycle the same story at different altitudes?

- **Pipeline altitude** — the tactical view: one pursuit moves through sense → qualify → score → pursue → design → audit → submit → deliver → learn.
- **GROWTH altitude** — the work-plan view: Gauge → Research → Organize → Write → Target → Health is the same pipeline told as *what work gets done this week*.
- **Lifecycle altitude** — the strategic view: the pursuit is an ORBITAL that either fails (back to dream) or wins and becomes a World.

The pipeline is the shape, GROWTH is the work plan, the lifecycle is the meaning — one story at three altitudes.

3. What is the cosmology loop, and why does it compound?

The cosmology loop is the self-reinforcing cycle in which a won-and-delivered World produces the inputs for the next Dream: **past performance** (evidence that wins the next pursuit), **relationships** (customers who trust you), and **surplus** (revenue that re-seeds the next pursuit). It compounds because each loop raises the next pWin: more wins → more past performance → better odds → more wins. A mature organization is running dozens of loops, and the Worlds are feeding the Dreams.

Doctrine 08 — Tools Change, Concepts Don't

1. What are the two layers of this course, and which one do we test on?

Layer one: the durable concepts — the market, the documents, the pipeline, the gates, the scoring, the ORBITAL, the lifecycle. **Layer two: the tools** — the software, portals, and AI systems that serve the concepts. We teach layer one and **only mention** layer two; **we test on the concepts**, never on a tool.

2. Name three durable facts from this doctrine that will still be true in twenty years.

Any three of: (1) Congress appropriates and agencies spend public, published money; (2) agencies will always publish their needs and rules in advance, and winners will be chosen on published criteria; (3) pursuits will always pass through some version of sense → … → learn with gates between; (4) every serious pursuit will always need a skeleton with an objective, resources, budget,

indicators, route, activities, and conditions; (5) ideas will always need structure before momentum, and momentum will compound into entities.

3. When a tool changes, what gets updated and what must not change?

When a tool changes, the **mapping between the doctrine and the tool** gets updated — quietly, in the alignment layer (the concept system map) — and the course goes on. What must **not** change is the **doctrine itself**: the concepts, the government facts, and the shared language between the human side and the machine side.

Part B — Module comprehension checks

Shared Core

S1. What is the difference between a contract and a grant? What do UEI, CAGE, and NAICS do?

A **contract** is the government buying goods or services from you (acquisition rules); a **grant** is the government funding a mission you carry out (grants policy). **UEI** is your 12-character identity in federal systems (free from SAM.gov); **CAGE** is your 5-character DoD/NATO identity; **NAICS** is the 6-digit code that places your business in an industry bucket and decides size standards and many set-asides.

S2. Name the nine stages in order. What are the four parts of a gate?

Sense, qualify, score, pursue, design, audit, submit, deliver, learn. A gate has a **trigger**, **criteria** (written in advance), a **decision** (proceed / proceed with conditions / stop), and an **owner**.

S3. What are the six questions every solicitation answers? What is the difference between a compliance reading and a strategy reading?

What is being bought/funded; who is eligible; how we will be judged (evaluation criteria — the most important); when the milestones are; how the response must be structured; where it must go and how. The **compliance reading** is the mechanical "what is required" pass that produces the compliance matrix; the **strategy reading** is the "what does the evaluator actually care about" pass. Winning proposals do both.

S4. What are the seven axes? What does "structure first, document second" mean?

Objective, Resources, Budget, Indicators, Transport, Activities, Logistics. "Structure first, document second" means the skeleton comes before the prose:

build the seven-axis structure before writing the proposal, because the proposal (and the score, and the gates) all trace back to the structure.

S5. What two numbers make a score? What is price-to-win?

pWin × value. Price-to-win is the price that is low enough to be competitive and high enough to be real — a position set at the intersection of your cost floor, the competitive range, and what the evaluation rewards.

S6. What moves a Dream to an ORBITAL, and an ORBITAL to a World? What is the cosmology loop?

Structure moves a Dream to an ORBITAL (the seven axes, a budget, an owner); winning and delivering moves an ORBITAL to a World. The cosmology loop is the self-reinforcing cycle where a delivered World produces past performance, relationships, and surplus that re-seed the next Dream — raising the next pWin, so it compounds.

Strategic Initiative track

SI-1. What are the two layers, and which is durable? Why is the doctrine the shared language between human and machine?

The durable concepts (layer one) and the swappable tools (layer two); the concepts are durable. The doctrine is the shared language because when a human says "gate," "pWin," or "ORBITAL," the machine must mean the same thing — the doctrine is what is encoded, so both sides can trust they are reasoning about the same concepts.

SI-2. What does a gate do to portfolio conversion? Why does the threshold discipline matter more than its exact value?

A gate is a conversion point in the funnel — it filters which opportunities advance, so it directly controls where attention and money are converted (or lost). The threshold's discipline matters more than its value because the discipline forces the honest "does this clear the bar?" conversation before resources are spent; the exact number is far less important than the fact that a bar exists and is enforced.

SI-3. What are the four parts of a gate? What is the difference between owner-with-standing and owner-with-process?

Trigger, criteria, decision, owner. **Owner-with-standing** has authority from position (a senior person — fast, decisive, only as good as their information). **Owner-with-process** has authority from the rule (e.g., the compliance officer who says "this cannot ship without the signature"). Healthy organizations mix both.

SI-4. Why is a firm's certification posture a strategy, not paperwork? What makes the STTR model structurally different?

Certifications decide *which pools a firm may enter* — HUBZone, 8(a), WOSB/VOSB each open a restricted, smaller competitive pool, so choosing and maintaining them is a go-to-market decision, not an administrative chore. STTR is structurally different because it **requires a nonprofit research institution partner** (a university or lab) and mandates work-share splits between the small business and the institution — it makes partnership a legal requirement, not an option.

SI-5. What is the difference between DREAM, GROWTH, and SuperNova as accents of one doctrine?

DREAM (Discover, Relate, Empower, Advise, Motivate) is the *engagement* accent — how you relate to the customer and mission. GROWTH (Gauge, Research, Organize, Write, Target, Health) is the *production* accent — the work arc of one pursuit. SuperNova is the *strategic* accent — the full arc from seed idea to operating entity. All three are the same pipeline told from different points of view.

SI-6. What does a healthy pipeline look like as a set of metrics? How do gates protect attention, money, and reputation?

A healthy pipeline shows conversion at each gate (how many sensed → qualified → scored → pursued → submitted), a portfolio value (sum of score × value), a win rate that rises over time, and a learning loop that demonstrably improves pWin factors. Gates protect **attention** (bid/no-bid stops pursuit-spend on weak fits), **money** (threshold stops spend below the bar), and **reputation** (readiness gates stop unready submissions that would embarrass the firm).

MBA track

MBA-1. What is the expected-value arithmetic of a pursuit? What is the sunk-cost trap?

Expected value = pWin × value; the expected return on pursuit spend compares that to the cost of pursuing (e.g., \$800,000 × 0.45 = \$360,000 expected value against a \$30,000 pursuit cost). The **sunk-cost trap** is continuing to spend on a pursuit because of what you have *already* spent — the rational test is forward-looking: *will additional spend clear the gate?* What is spent is spent.

MBA-2. What are the three inputs to price-to-win? What is the difference between LPTA and best-value?

(1) What it costs us (the floor); (2) what the market will bear (the competitive range); (3) what the evaluation rewards (price posture). **LPTA** (lowest price technically acceptable): once the technical bar is passed, price nearly decides it — so price low to win. **Best value**: the agency weighs technical strength against price — so a defensible higher price with a stronger technical story can win.

MBA-3. How does the ORBITAL work as a business plan? What is the difference between the budget axis and the indicators axis?

The ORBITAL is a business plan in miniature: Objective (the value proposition), Resources (the team), Budget (the cost structure and revenue), Indicators (the measurement layer), Transport (route to customer), Activities (the plan), Logistics (the conditions). The **budget axis** is *what it costs and what it's worth*; the **indicators axis** is *how you know you're on track*. Budget is the business test; indicators are the steering wheel.

MBA-4. How does NAICS determine a firm's competitive pool? What does HUBZone status do for a firm's market position?

NAICS assigns the industry bucket, which sets the **size standard** and the **eligibility** for set-asides — so the NAICS code on a solicitation determines who is allowed to compete. HUBZone status moves a firm into a **restricted pool** (set-asides and evaluation preferences reserved for HUBZone-certified firms) — a smaller, friendlier competitive field, which is a real market advantage.

MBA-5. What is the cosmology loop, and why does it compound? What role does past performance play?

The cosmology loop is the self-reinforcing cycle where a won-and-delivered World produces past performance, relationships, and surplus that seed the next Dream. It compounds because past performance is the evidence evaluators weight heavily — each delivered win raises the next pWin, and higher pWin means more wins. Past performance is the compounding asset that turns a firm from a newcomer into a trusted supplier.

MBA-6. What makes a pursuit portfolio defensible? How does price-to-win connect to the business model?

A defensible portfolio is a deliberate allocation of scarce capacity across risk levels — safe wins, mid bets, long shots — that clears the threshold discipline and can be defended to a board with numbers. Price-to-win connects to the business model because it is where the model's margin is decided: a price above the floor preserves margin; a price that wins at any cost produces a business that loses money on the work it wins.

MPA track

MPA-1. What are the steps of the money flow? What is the difference between contract and assistance dollars?

Congress appropriates → the budget process allocates → agencies turn appropriations into programs and purchases → solicitations publish the need → pursuers compete → winners perform → payment flows → awards are published → published history feeds the next cycle. **Contract dollars** buy goods/services for the government; **assistance dollars** (grants and cooperative agreements) fund

missions the government shares with a recipient. Different rulebooks, different portals, same published flow.

MPA-2. What is the FAR, and what does it protect? How do grants policy and acquisition rules differ?

The **FAR** (Federal Acquisition Regulation) is the rulebook for how the government buys things — the mechanism that makes spending fair, open, and competitive. **Acquisition rules** govern contracts (buying goods/services); **grants policy** governs assistance (funding missions). Both protect fairness and accountability, but they are separate rulebooks with separate vocabularies and separate portals — the RFP for contracts, the NOFO for grants.

MPA-3. What are the post-award obligations of a grantee? Why must budget and work tell one story?

Post-award obligations include performing the funded work, **reporting** (progress, financial, and final reports), spending within the approved budget and the agency's cost rules, and complying with the award's terms (including, for STTR-style awards, patent-rights reporting). Budget and work must tell **one story** because evaluators and auditors check alignment: an inspiring narrative with a disconnected budget signals the applicant is not serious about spending public money well — and it is a reviewed criterion.

MPA-4. What is the statutory purpose of small-business set-asides? How is eligibility an equity instrument?

The statutory purpose is to ensure a portion of federal contracting reaches small businesses — a deliberate policy correction so the federal market is not a private club. Eligibility is an equity instrument because it **narrows the pool by statute**: HUBZone, 8(a), WOSB/VOSB each reserve competitions for firms that meet a defined disadvantage or underutilized-zone test, giving those firms a seat at a table they would otherwise be priced out of.

MPA-5. What is the oversight chain for federal money? How is transparency a governance mechanism?

The oversight chain runs from **reporting** by the awardee, through **agency review**, **audits** (inspectors general, GAO), and finally the **published record** — with the public as the final reader. Transparency is a governance mechanism because **published awards let anyone trace the money**: they make spending accountable to citizens, enable external scrutiny, and supply the past-performance evidence the next competition runs on.

MPA-6. How are the money flow, the rules, the programs, and oversight one system? What is the equity lens on the doctrine?

They are one system because they are the same transaction viewed at different layers: the money flow is *how* the dollar travels, the rules are *what keeps the travel fair*, the programs are *where Congress deliberately steered the money* (small-business set-asides), and oversight is *how the whole thing is kept honest.*

The equity lens asks who the system serves and who it leaves out — whether the small-business programs actually deliver their intended benefit, and whether public money is spent in a way that corrects rather than compounds disadvantage.

Part C — Practice bank answers

Part A — Doctrine drills

D01-E1. The three kinds: **contracts** (government buys goods/services), **grants** (government funds a mission you carry out), and **other transactions / cooperative agreements** (in-between, faster shapes for prototypes and partnered research).

D01-E2. The problem is **not** the identifiers — Ravonics has a valid UEI, CAGE, and NAICS set; the problem is a **capability fit**. The two checks a smart capture lead runs: (1) *Do we have the capability and past performance to perform this?* (No — community health is outside the four NAICS and the bench.) (2) *Is this in our strategy/portfolio?* (No.) Declining is the correct professional outcome.

D01-E3. It **cannot** compete until the SAM registration is current — the registration is the eligibility precondition, and an inactive record is a disqualifying gap regardless of technical strength or certification. The first thing to verify is the **SAM registration status** (and that the UEI/CAGE are current), because eligibility is a hard gate: *if you are not registered, nothing else matters.* This is why doctrine/01's memorization sentence — "you register once in SAM.gov… every pursuit starts with 'is our SAM record active?' " — is the practical rule.

D02-E1. The six questions: what is being bought/funded; who is eligible; how we will be judged; when the milestones are; how the response must be structured; where it must go and how. The **evaluation criteria** ("how we will be judged") is what a winning proposal gets written to — a proposal is a document written to the evaluation criteria.

D02-E2. (a) Page limits and format rules → **Section L** (instructions); (b) the submission portal → **Section L** (submission mechanics); (c) relative weights → **Section M** (evaluation factors). Section L is the rules of the game; Section M is the rubric.

D02-E3. The lead is right. The visible deadlines are the questions deadline and the submission deadline; the invisible third deadline is the **internal deadline**. Submitting at 4:30 PM Friday leaves no room for a file-validation failure, a portal jam, or a last-minute correction — and "well, it was 4:59:58" is not a legal argument. The internal deadline (Wednesday) builds in the review, fix, and mechanical-submission buffer that the readiness gate requires.

D03-E1. Sense → qualify → score → pursue → design → audit → submit → deliver → learn (one-line jobs in Part A, Doctrine 03, Q1 above).

D03-E2. (a) **Sense** — reading a forecast is watching the river. (b) **Audit** — reviewing a draft against the evaluation criteria is the color-team discipline. (c) **Deliver** — shipping a milestone is delivery. (d) **Learn** — a win/loss meeting that updates pWin factors is the learning loop.

D03-E3. Pursuing all three at once violates the pipeline's core discipline: capture capacity is finite, and a firm that pursues everything goes bankrupt. The gate exists to force scarcity into the open. The alternative: **qualify and score all three**, set a threshold, and *choose* the one (or two) that clear it — a deliberate portfolio decision instead of a volume bet. A disciplined "no" to two of them is a professional outcome, not a failure.

D04-E1. The four parts: trigger, criteria, decision, owner. The missing part that turns a gate into a rubber stamp is **criteria written in advance** — when the criteria are invented at the meeting, the gate is theater.

D04-E2. A defensible bid/no-bid gate for a five-person Ravonics: (1) *Are we eligible?* — evidence: SAM active, certification/code matches the set-aside. (2) *Is it real and fitting?* — evidence: a real solicitation whose scope is inside our four NAICS. (3) *Does the value clear our floor?* — evidence: published award size our $X floor. (4) *Do we have capacity?* — evidence: this would be our one active pursuit; staffing named. (5) *Can we comply?* — evidence: every compliance item has an owner. Each criterion is answerable yes/no and the evidence is stated — that is what makes it a gate rather than a mood.

D04-E3. The **owner** decides — but the resolution is process, not personality. Because the criteria are written in advance, the disagreement is resolved against the criteria, not against the founder: the compliance officer assembles the evidence that the pursuit fails a published criterion; the founder (owner-with-standing) makes the final call, but must answer to the criterion and the evidence. If the founder overrides, the override is recorded and the cost is visible. The point is that the *decision is boring* — it is a check of boxes, not a clash of wills. Healthy organizations use both owners: the process owner keeps the bar honest; the standing owner takes accountability for the final call.

D05-E1. The score is **pWin × value**; multiplying matters because it forces both likelihood and worth into one comparable number. Price-to-win is the price that is low enough to be competitive and high enough to be real — a position, not a guess.

D05-E2. (a) $\$200{,}000 \times 0.60 = \mathbf{\$120{,}000}$. (b) $\$2{,}000{,}000 \times 0.05 = \mathbf{\$100{,}000}$. (a) scores higher despite being ten times smaller — the score says so without an argument. Against a $120,000 threshold: (a) is a **marginal pass / BID** (at the line); (b) is a **NO BID** (below). Note the threshold's value matters less than the discipline: the point is the conversation happens before resources are spent.

D05-E3. A defensible price is a **position**: stay above the $95 floor, inside the $105–$125 range, and lean toward the **upper part of the range** because (a) the incumbent is well-liked, so a price war you can't win is suicide; (b) best-value rewards a stronger technical story, so price slightly above the competitive midpoint can be justified by capability. Recommendation: ~$118–$122/hour with a technical story that earns the premium — not the floor (which signals you can't justify value) and not the ceiling (which invites a sharp-pencil rejection). A too-low price would win a contract at a margin that bleeds the firm — the worst outcome.

D05-E4. A model rubric for a 541330 engineering-services set-aside: factors (weights) — customer relationship/incumbency (0.25), past performance (0.25), technical approach (0.20), price competitiveness (0.15), capacity/capture readiness (0.15). Score each 0–1 with an evidence column; weighted pWin = sum(weight × score). Composite = pWin × value. Set a threshold (e.g., pWin 0.42 AND composite 0.42 × value) and make the call. Grading check: the answer is defensible if every score has stated evidence and the call follows from the stated threshold. (See the worked exemplar's Step 2 for a fully worked version.)

D06-E1. Objective (what are we trying to do, what counts as success), Resources (who is on the team), Budget (what does it cost, what is it worth), Indicators (how do we know we're on track), Transport (how do the pieces reach the customer), Activities (what work, in what order), Logistics (what must be true for the work to happen).

D06-E2. The **Resources** axis is empty (no named team) and the **Activities** axis is empty (no plan of work). The ORBITAL is a wish rather than a plan because the objective and money cannot be executed: without a team nothing is staffed, and without activities nothing is scheduled. An ORBITAL with no team and no plan cannot be delivered — and the pursuit would disintegrate the moment the capture lead got busy.

D06-E3. A build exercise — no single right answer. Grading check: all seven axes filled with *specific, sourced* content (from the solicitation and the market), budget and indicators defensible, and the "hardest axis" answer honest about what it reveals (a hard Logistics axis usually reveals a clearance, subaward, or facility gap; a hard Indicators axis usually reveals a fuzzy objective). Model: the worked exemplar's Step 4.

D07-E1. Dream (seed idea) → ORBITAL (structured pursuit) → World (fully realized enterprise). Structure moves Dream → ORBITAL; winning and delivering moves ORBITAL → World.

D07-E2. Ravonics is an **ORBITAL-stage firm**: it has structure (registrations, certifications, a partner, a bench) and its first structured pursuit is under way — but it is not yet a World (no delivered, revenue-generating contract that produces past performance). The two moves that advance it: (1) **win and deliver its first ORBITAL** (the STTR Phase I) so it gains past performance; (2) **convert that delivery into the next pursuit's evidence** — feeding the

cosmology loop. Its HUBZone status and INSTAR partnership raise its pWin in the pools those programs open, which is precisely the kind of deliberate bet that moves an ORBITAL toward World-hood.

D07-E3. The firm is fragile because it has **stopped feeding the loop**. The cosmology loop compounds only when Worlds produce past performance, relationships, and surplus that re-seed the next Dream. A firm that wins one contract and stops sensing and qualifying will, when that contract ends, have no pipeline behind it — the single contract is a one-time event, not a compounding asset. The profitable present is masking an empty future; the pipeline must keep turning.

D08-E1. The two layers: durable concepts (layer one) and swappable tools (layer two); we test on the concepts. When a tool changes, the **mapping** (alignment layer) gets updated, and the **doctrine never changes**.

D08-E2. None of those changes touch the doctrine. A new portal interface, a vendor's template change, and an AI drafting first pass are all **tool-layer** changes — they serve the same concepts (the market, the pipeline, the gates). They belong in the **alignment layer** (the concept system map), not in the doctrine files. The concepts — solicitations publish needs and rules in advance, winners are chosen on published criteria, pursuits pass through the pipeline — are untouched.

D08-E3. A strong two-sentence answer: *"I haven't used your specific system, but I know the doctrine it holds — the pipeline, the gates, the ORBITAL structure, the scoring discipline. Show me your system and I'll map the concepts onto it quickly."* That demonstrates the durable layer (portable) versus the tool layer (learnable) — which is the entire point of doctrine/08.

Part B — Shared-core practice

SC-E1. A build/apply exercise. Grading check: the map has three labeled steps — an appropriation (the law/amount), a solicitation (the announcement), and a past award (the published winner + amount) — traced on USAspending.gov, and it notes what the record reveals. The worked exemplar's money context (Step 1) and the literacy money-flow map are the models.

SC-E2. Apply. The answer is defensible if it places the company's activity on the pipeline and names a gate with criteria. Example shape: "The company sensed a forecast, qualified, scored, pursued, and won — the decisive gate was the threshold gate, where the pursuit cleared the pWin/value bar, or the readiness gate, where a compliant submission was required." The pipeline vocabulary must be visible.

SC-E3. Build. Grading check: at least eight rows, each a real "shall" with the section it came from and where the response satisfies it. Model: the worked exemplar's Step 3 compliance matrix.

SC-E4. Apply. The strategy reading must quote the evaluation weights and adjectives and state what story the proposal must tell to maximize them. Grading check: the reading mirrors the criteria in the order the agency weighted them and names the biggest threat to the score.

SC-E5. Build. Grading check: pWin argued factor by factor with evidence; value grounded in the solicitation's numbers; composite = pWin × value; a threshold stated in advance; a one-paragraph BID / BID WITH CONDITIONS / NO BID call. Model: the worked exemplar's Steps 2 and 5.

SC-E6. Apply. See D07-E2 for the model: Ravonics is ORBITAL-stage; HUBZone + STTR partnership raise pWin in restricted pools; the moves that advance it are winning/delivering the first ORBITAL and converting delivery into past performance.

Part C — Track practice

SI-E1. Build. Grading check: the doctrine statement names durable concepts (market is published, solicitations answer six questions, the pipeline, gates with criteria, the ORBITAL, the lifecycle) written so a new hire could apply them **without knowing the current software**. A doctrine statement that mentions a product name fails the stack-agnostic test.

SI-E2. Build (portfolio). Grading check: a fixed capacity budget; a ranking with gate reasoning (eligibility → score → capacity); an honest rejection of the pursuit that is eligible but unstaffable (the IDIQ) — a disciplined "no" is the professional outcome; the recommendation is defensible to a board. Model reasoning: the STTR topic fits capability + certification; the HUBZone set-aside fits certification but may pay less; the IDIQ is a doorway you cannot staff — decline or defer.

SI-E3. Build (gate charter). Grading check: for each gate (bid/no-bid, threshold, review, readiness) the charter states the **trigger, criteria written in advance, evidence, and decision owner**. Model: the worked exemplar's Step 5 shows the threshold gate as a written-in-advance rule.

MBA-E1. Expected value = \$800,000 × 0.45 = **\$360,000**; expected return on pursuit spend = \$360,000 − \$30,000 = **\$330,000 net expected value**, or ~11× return on pursuit spend. It is rational to continue while the expected value of additional spend exceeds its cost. The sunk-cost test is *forward-looking*: at what pWin does it stop being rational? Solve \$800,000 × pWin = \$30,000 → pWin **0.0375** — below roughly 4%, the pursuit is no longer worth pursuing regardless of what has already been spent.

MBA-E2. Build (price-to-win). Grading check: the three inputs are shown (cost floor, competitive range, evaluation posture); the recommended price sits inside the range and above the floor; the one-paragraph story justifies the price as a position; and the LPTA vs best-value distinction visibly changed the price

(LPTA → price at/near floor to win; best-value → price with headroom for a technical story). Model: D05-E3 above.

MBA-E3. Build (flywheel). The arithmetic should show: Year 1 portfolio value = pWin × value × number of pursuits; after each delivered win, past-performance evidence lifts the pWin factor (e.g., +0.10 on the past-performance factor), so pWin > pWin , so Year 2 portfolio value > Year 1 — and so on. Grading check: the compounding is explicit (each delivered win raises the next pWin) and the numbers are internally consistent.

MPA-E1. Apply. Grading check: the answer traces a published award back to an agency and appropriation context, states what the record reveals (agency, amount, often the contractor and purpose), and names what it does **not** show (the full story behind the award — the qualitative evaluation, the losing proposals). Model frame: the money-flow map in literacy/where-the-money-flows.md.

MPA-E2. Apply. Grading check: the answer finds the compliance requirements in each document (the "shall"s, certifications, clauses), and compares how each enforces fairness — the RFP via acquisition rules and Section L/M format + evaluation commitments; the NOFO via grants policy and review criteria + budget alignment. Both protect fairness; different rulebooks.

MPA-E3. Decide (policy memo). Grading check: the memo states the program's statutory purpose, uses Ravonics as the concrete object, and argues whether the program achieves its intent with evidence. A defensible answer engages the equity question: set-asides narrow the pool by statute to correct under-participation — but their success should be measured by whether eligible firms actually win and deliver, not by how many apply. Model: HUBZone exists to steer work to underutilized zones; Ravonics's HUBZone status is a real, structural advantage in restricted pools — the policy's intent and the firm's strategy align.

Part D — Integrated practice

INT-E1. Build (mini-capstone). Grading check — the five artifacts are all present and consistent: compliance matrix, strategy reading, ORBITAL (seven axes), score (pWin argued factor by factor + value + composite + threshold), gate charter, and a written BID / BID WITH CONDITIONS / NO BID call. The worked exemplar is a complete model answer; the capstone rubric is the grading contract.

INT-E2. Decide (post-mortem). Grading check: the post-mortem looks up the actual published award; compares the outcome to the score's prediction; updates at least one pWin factor with what the outcome teaches; and states honestly what the score got right and what it missed. The learning loop (doctrine/03, stage 9) is the point — this is the move that turns a practice exercise into a career skill.