

Sample Capstone — From NOFO to Submission-Ready Pursuit

A complete, worked undergraduate capstone: a small HUBZone firm takes a fictional federal services recompile and runs it through the full pursuit pipeline in 14 weeks — discovery, score, capture, gates, win strategy, four-volume outline, compliance matrix, color team, orals briefing, and submission readiness. Every step shows the reasoning, not just the result.

What this is. The assignment brief for the undergraduate capstone lives in `course/undergraduate-480-capstone-build-an-orbital.md` (and its shared base, `course/capstone-build-an-orbital.md`). This page is a **model answer** — a sample project that shows what an excellent submission looks like, milestone by milestone, with the grading rubric the instructor uses at every step. Read the brief and this sample together. This sample is deliberately built on a **fictional agency and requirement** so it can be worked freely in a classroom; a student's own capstone uses a real, published, closed solicitation, unless the instructor assigns this fictional sample as the teaching object.

The stack-agnostic rule applies throughout (doctrine/08): no deliverable below requires any product or system. Your organization's tooling — whatever it is — is how you produce the artifacts; the concepts are what is graded.

Tier: undergraduate (B.S. in Government Business Development) · **Course anchor:** DL 480 · **Term:** 14 weeks · **Teams:** 3–4, each member owning at least one deliverable · **Defense:** individual, before a skeptical panel · **Tools:** none required beyond paper, a whiteboard, or a spreadsheet.

The sample scenario

The fictional agency and requirement

FEDERAL RESILIENCE AND ADAPTATION SERVICE (FRAS) — a *fictional* independent federal agency responsible for the resilience of critical infrastructure against climate and geospatial hazard. FRAS is recompeting its analytic support contract.

FRAS RFP 2026-014 — Enterprise Geospatial and Infrastructure Analytics Services (EGIAS). The successor contract for FRAS's coastal and riverine infrastructure analytics mission. The contractor will provide geospatial data management, infrastructure condition assessment, hazard modeling, and decision-support analytics to FRAS's Resilience Operations Directorate.

Key solicitation facts (abridged for teaching):

Fact	Value
Vehicle	Single-award services contract (FAR Part 15), HUBZone set-aside
Period	1 base year + 4 option years
Value	\$6.5M all-options (base year \$1.3M)
Incumbent	Meridian Analytics Group (fictional) — 5-year incumbent, well-liked by the program office
Evaluation	Best-value tradeoff: Technical (45%) · Management (25%) · Past Performance (15%) · Price (15%)
Submission	Electronic portal; 100-page total proposal limit (all volumes); Section L format rules
Key requirement	The contractor must maintain a named Program Manager and at least two senior analysts for the full period of performance

This is a **recompete**, not a new-start — which is exactly why it is a rich undergraduate capstone. The incumbent is strong, the customer knows the work, and the winning strategy must be argued against a known competitor, not against a vacuum. (The same 14-week plan works for a grants-style NOFO/SBIR pursuit with the budget and technical volumes swapped in; the literacy map `literacy/how-to-read-a-nofo.md` shows where the two documents differ.)

The practitioner firm: Ravonics LLC

The team practices as **Ravonics LLC** — the canary company from the case study. A HUBZone-certified small business in Ravenswood, West Virginia. UEI YCBDVKN1A9G7, CAGE 20D58. Four NAICS codes: 541512 (computer systems design), 541330 (engineering services), 541715 (R&D in physical and engineering sciences), 518210 (data processing and hosting). A bench of PhD physicists and AI specialists. One research partner: **INSTAR Lab**, the nonprofit research institution that makes Ravonics STTR-eligible. Ravonics has never held a FRAS contract, but its geospatial modeling and data-processing capabilities map cleanly onto the EGIAS statement of work — and it is **HUBZone-certified**, which is the eligibility ticket into this set-aside.

The problem statement

The team’s assignment. Produce, in 14 weeks, a submission-ready pursuit package for the EGIAS recompetete: a score and bid/no-bid call argued factor by factor and defended against the *criteria you write in Week 3*; a complete capture plan; a seven-axis ORBITAL; a full four-volume outline with a compliance matrix; a data-grounded price-to-win position; a color-team review log; an orals briefing; and a readiness call. Then present the package to a skeptical panel and — after the competition would have run — do the post-mortem against the published award.

The 14-week milestone plan, mapped to the pipeline

Each week is a pipeline stage doing its one-line job (doctrine/03), with a gate where the doctrine requires one (doctrine/04). The default threshold for the score gate is the pursuit pipeline’s canonical **0.42** — the bar the ecosystem’s Dream pursuit pipeline uses for “worth pursuing.”

Week	Pipeline stage	Milestone	Deliverable
1	sense	Kickoff: read the EGIAS solicitation as a document	Team formed; the 15-minute scan; solicitation fact sheet
2	sense → qualify	Discovery and eligibility	Eligibility memo + bid/no-bid gate (first gate)
3	score	Score the pursuit	Score: $pWin \times value$, threshold 0.42 (threshold gate)
4	pursue	Capture plan — knowing the customer and the fight	Capture plan draft
5	pursue	Win strategy + gate charter + ORBITAL	Gate charter (graded first) ; win-themes sheet; ORBITAL draft
6	design	The shape of the answer	Four-volume outline + compliance matrix (first pass)

Week	Pipeline stage	Milestone	Deliverable
7	design	Volume I — Technical	Technical volume draft (storyboarded)
8	design	Volume II — Management + Volume III — Past Performance	Management + Past Performance drafts
9	design	Volume IV — Pricing + price-to-win	Pricing volume draft + price-to-win position
10	audit	Color team — Pink then Red	Color-team findings log (independent reviewers)
11	audit → readiness	Revision + orals briefing	Revised package; orals briefing outline
12	submit	Submission readiness	Readiness gate: white-glove check, compliance re-run, final package
13	(submit)	Dry-run of the presentation	Presentation rehearsal
14	learn	Defense + post-mortem	Final package + honest post-mortem vs. the published award

Per-milestone deliverables — what “done” looks like

Week 1 — sense: read the document

Deliverable: solicitation fact sheet + 15-minute scan. The team reads EGIAS as the agency wrote it, not as they hope it is. The fact sheet answers the six questions every solicitation answers (doctrine/02): what is being bought, who may bid, what the agency will weigh, what must be submitted, when, and where. The scan names the *three deadlines* — questions deadline, submission deadline, and the team’s internal deadline (set to leave a full week for review and mechanical submission).

Week 2 — sense → qualify: the bid/no-bid gate

Deliverable: eligibility memo + a gate decision. The first gate asks four yes/no questions, answered with evidence: *Is FRAS a customer worth knowing? Do we have the past performance? Does the set-aside fit us — are we eligible and are we actually competitive in a HUBZone pool? Do we have the capacity?* The memo states the call — **BID** — and the evidence. A disciplined “no” would have been a professional outcome too; the team argues *why* this one is a yes.

Week 3 — score: the threshold gate (0.42)

Deliverable: scored pursuit sheet. The team scores pWin **factor by factor** against the seven-factor composite (past performance, technical capability, price competitiveness, incumbent/relationship advantage, team strength, customer knowledge, strategic alignment — doctrine/05), with **evidence for every factor** and at least two factors citing public data. Value is grounded: the estimated contract value and the pursuit cost. The composite: **pWin 0.52 × value \$6.5M → score clears the 0.42 threshold.** The gate is recorded: **PURSUE**, with conditions (the incumbent advantage is the known weakness; the win strategy must answer it).

Week 4 — pursue: the capture plan

Deliverable: capture plan. The capture plan is the strategy for winning *this* competition: the **customer map** (who at FRAS decides, who influences, what they are measured on), the **competitive landscape** (the incumbent’s likely position, who else can meet a HUBZone set-aside), the **incumbent analysis** (why Meridian has held the work, where it is vulnerable), and the **discriminator plan** (what Ravonics will be *demonstrably better at*, not just different at). The capture plan names the ghost themes the team will suppress (doctrine/09) — the things Ravonics could claim but must not, because they would not survive Red Team.

Week 5 — pursue: win strategy, gate charter, and ORBITAL

Deliverable (graded first): the gate charter. Before building anything else, the team writes the criteria for every gate the pursuit will cross — bid/no-bid, threshold, review (color-team exit criteria), readiness — in advance, exactly as doctrine/04 demands. The final score and call in Week 12 are judged against *these* criteria, not the ones the team would write after knowing the answer. **Deliverable: win-themes sheet.** The **win themes** (the reasons an evaluator will rate us higher) and the **discriminators** (the things only we can credibly claim): a named anchor partnership with INSTAR Lab, a quantified outcome target for hazard-model accuracy, a sustainability mechanism for the data assets the contract will leave behind. **Deliverable: ORBITAL draft.** All seven axes — Objective, Resources, Budget, Indicators, Transport, Activities, Logistics — filled with real, sourced content (doctrine/06). The transport and

logistics axes carry the small-town delivery reality: a Ravenswood-based team delivering national work.

Week 6 — design: the shape of the answer

Deliverable: four-volume outline + compliance matrix. The **compliance matrix** maps every “shall” in EGIAS to the section and page that will answer it — the internal spine of the whole effort. The **four-volume outline** allocates the page budget to the factors that carry the weight (doctrine/05, the literacy map): the 45% Technical factor gets the best pages and the deepest storyboarding; the 15% Price factor gets a clean, defensible cost buildup, not fluff.

Week 7 — design: Volume I — Technical

Deliverable: Technical volume draft. Understanding of the requirement (restated with insight, never parroting the SOW), the technical approach answering every technical requirement, risk management, quality control, security, and the transition plan that answers the incumbent’s biggest weapon — the cost and friction of switching. Every page asks the evaluator’s question: *does this show real understanding, and can this team actually deliver?*

Week 8 — design: Volume II — Management + Volume III — Past Performance

Deliverable: Management and Past Performance drafts. The Management volume shows who runs the work and how: the named **Program Manager** the RFP requires, the org chart, staffing, key personnel, subcontract management (INSTAR Lab’s role), and schedule management. The Past Performance volume is the **evidence selection discipline**: three to five recent contracts, prioritized by relevance to EGIAS, with the header, relevance, performance narrative, and reference for each. The team deliberately excludes anything with pending claims or unresolved performance issues — one bad citation can sink the volume.

Week 9 — design: Volume IV — Pricing + price-to-win

Deliverable: pricing volume + price-to-win position. The cost buildup (labor categories and hours, other direct costs, travel, subcontract, indirect rates and fee), organized by CLIN and with a **basis of estimate** for the largest line items. The **price-to-win** frame: the cost floor, the estimated competitive range, and a recommended price argued as a *position* against the evaluation posture. Because FRAS is buying best-value on a 45/25/15/15 split, the price is set to be *credible and competitive*, not lowest — and the volume must reconcile with the Technical volume line for line.

Week 10 — audit: Pink then Red

Deliverable: color-team findings log. The class runs the review sequence with the facilitation discipline (literacy/how-color-teams-work.md): **reviewers independent of writers, a locked roster, specific findings (factor + evidence + revision), forward-only cadence.** **Pink** checks compliance and coverage against Section L/M. **Red** scores the full draft as a government evaluator, factor by factor, and reconciles the technical and cost stories. The findings log is the graded artifact: every finding names the factor at issue, the missing or weak evidence, and the specific revision. The discipline's blunt rule applies: **any Red rating on a critical factor blocks submission.**

Week 11 — audit → readiness: revision + orals briefing

Deliverable: revised package + orals briefing outline. The Red findings are resolved in the revision. The **orals briefing** prepares the team for the post-submission conversation (doctrine/10): the story in one line, the discriminators, the answers to the questions a panel will ask about the incumbent advantage and the transition risk. The briefing is the presentation skeleton — the defense in Week 14 uses it.

Week 12 — submit: the readiness gate

Deliverable: final package + readiness call. The **readiness gate** is the last gate that can still stop a bad submission. The white-glove pass (formatting, forms, signatures, page counts, portal mechanics) and a full compliance re-run are checked against the gate charter written in Week 5. The team records the readiness decision: **READY**, with the residual risks named (the incumbent's incumbency is the known risk; the mitigations are in the transition plan).

Week 13 — submit: the dry-run

Deliverable: rehearsal. The team presents the full package to a mock panel and fixes what the mock panel finds. The dry-run is itself a white-glove for the defense.

Week 14 — learn: defense and post-mortem

Deliverable: final package + post-mortem. The team presents (15 minutes) and each member defends their own slice against a skeptical panel. Then — because the competition would have run — the team looks up the **actual published award** (who won, for how much, on the real FRAS award or the nearest published analogue), compares it honestly to their pWin factors, and writes the half-page post-mortem: what the score got right, what it missed, and what the team would change next time. This is the learning loop (doctrine/03, stage 9) made visible — the moment the degree becomes a career skill.

Grading rubric per milestone

Every milestone is graded on the shared four-level scale — **Not Ready** · **Developing** · **Proficient** · **Distinguished** (see the capstone README for the full scale and the six shared dimensions). The weights below reflect the undergraduate tier's priorities: the ORBITAL and the score carry the most weight, exactly as the degree's assessment contract says.

Week / milestone	Criteria	Not Ready	Proficient	Distinguished
1–2 · Sense + bid/no-bid gate	Reads the document as the agency wrote it; eligibility determined; gate call argued	Misreads key facts; gate absent or rubber-stamped	Six-question fact sheet correct; four gate criteria answered with evidence	Names the third (internal) deadline; argues <i>why</i> BID is right, not just that it is
3 · Score 0.42	pWin argued factor by factor with evidence; value grounded; threshold call defended	Score asserted, not argued; no sources	Every factor scored with cited evidence; composite computed; gate recorded	Data-weighted factors; the incumbent factor named as the honest weakness
4 · Capture plan	Customer map, landscape, incumbent analysis, discriminators	Landscape guessed; incumbent analysis missing	Customer, competitors, and discriminators argued from evidence	Ghost themes named and suppressed; the plan would change if Meridian's posture changed
5 · Gate charter + win strategy + ORBITAL	Criteria written before the work; all seven ORBITAL axes sourced	Axes thin or unfilled; criteria after-the-fact	All seven axes filled and sourced; gate charter complete	Transport/logistics carry the small-town reality; the charter binds the team's later call

Week / milestone	Criteria	Not Ready	Proficient	Distinguished
6 • Four-volume outline + compliance matrix	Every “shall” mapped; page budget follows the weights	Partial coverage; page budget ignores Section M	Full compliance matrix; volume outline allocates to the 45% Technical factor	The matrix and the outline are cross-checked — nothing has a home and a home has nothing
7–8 • Technical + Management + Past Performance	The four volumes tell one story; evidence discipline	Volumes contradict each other; unsupported claims	Volumes consistent; named PM; PP citations relevant and clean	Transition plan answers the incumbent’s switching-cost weapon
9 • Pricing + price-to-win	Cost buildup real; BOE present; price a position	Price is a guess; BOE absent	BOE for major lines; price-to-win argued against competitive range	Price reconciles with Technical line-for-line; posture fits best-value
10 • Color team (Pink + Red)	Independent reviewers; specific findings; forward-only	Friendly review; vague findings	Findings name factor + evidence + revision; exit criteria met	Red ratings resolved before submission; the log shows the package improving
11–13 • Orals briefing + readiness	Briefing is the story; readiness gate honestly run	No briefing; readiness a formality	Readiness checked against the Week-5 charter	Residual risks named with mitigations before the panel finds them
14 • Defense + post-mortem	Survives the panel; learns honestly	Evasive; no post-mortem	Clear, credible, doctrine used under questioning	Post-mortem updates the pWin factors with what the actual award taught

Composite weight: Solicitation mastery 20 · ORBITAL quality 25 · Score and bid/no-bid 20 · Governance and gates 15 · Presentation and defense 20 · (*Post-mortem bonus +0 to +10, capped at 100*) — per the DL 480 brief. The panel’s defense grade is the final gate: it can move the composite by one level in either direction.

Where this maps to Shipley + Dream DNA

The industry canon. The EGIAS pursuit maps one-to-one onto the **Shipley business-development lifecycle** — *marketing* → *capture* → *proposal* → *negotiate* → *award* (*literacy/industry-standard-canon.md*). Weeks 1–3 are **marketing** (market segmentation, opportunity identification, the pursuit decision — the bid/no-bid and threshold gates are Shipley’s gates 2 and 3). Weeks 4–5 are **capture** (capture readiness — gate 4): the capture plan, the win themes, the discriminators, the gate charter. Weeks 6–9 are **proposal** (proposal planning and production — gate 5): the four-volume outline, the compliance matrix, the color-team cadence. Week 12 is **submission mechanics** (gate 6): the readiness gate and white-glove. Week 14 and the post-mortem are **negotiate and award** (gate 7): the orals briefing and the debrief/lessons-learned loop.

The Dream DNA. The same pursuit is the *machine side’s* object. Every artifact this team produces has a machine-corpus counterpart in the **DNA corpus** — the versioned package of rules, playbooks, and formulas that encodes the doctrine (*align/concept-to-system-map.md*): the **seven-factor composite** that this team’s Week-3 score operationalizes by hand, the **pWin × value** formula with the **0.42 threshold gate**, the **ORBITAL’s seven axes** the team fills in Week 5, the **four-volume bundle** shape the team outlines in Week 6, and the **golden-thread verification** that walks this exact pipeline — the same stages this 14-week plan walks — against the **canary company**, which is the same Ravonics this team is practicing as. When the human team builds the ORBITAL and the machine scores the pursuit, both are doing the Week-5 and Week-3 work on the same company through the same doctrine. That is what “one doctrine, two renderings” means — and it is why the capstone that makes this team fluent in the human rendering also makes them able to read, steer, and challenge the machine rendering when they meet it in a real capture organization.

The capstone in one sentence

Four years of doctrine and documents converge in one package: a real recompute, a defensible score that clears the 0.42 gate, a capture plan that answers the incumbent, an ORBITAL with all seven axes real, a four-volume outline with a compliance spine, a price-to-win that is a position not a guess, a color-team log that shows the package improving, a readiness call bound by a charter written in Week 5 — and the honest post-mortem that turns the degree into a career skill.