

Capstone Projects — Sample Program & Assessment Framework

The capstone is where the doctrine stops being read and becomes done. It is the final assessment of every tier — the moment a learner takes a real pursuit, runs it through the pipeline, and defends the result before people who are paid to be skeptical. This folder is the sample program: how to run a capstone, how to grade it, and a worked sample capstone for every rung of the ladder.

What this folder is, and is not. The assignment briefs for the degree tiers live in `course/`: the shared capstone (`course/capstone-build-an-orbital.md`), the undergraduate extension (`course/undergraduate-480-capstone-build-an-orbital.md`), the graduate culminating experience (GC 690 / GC 690T, `course/graduate-program-overview.md`), and the doctoral dissertation sequence (DL 903 / DL 990, `course/doctoral-program-overview.md`). The files in *this* folder are **sample projects** — complete, worked model answers and program-level guidance — so students can see the shape of excellence and instructors can see how to run and grade a capstone program across all three tiers. Read the brief and the sample together: the brief says what to do, the sample shows what done looks like.

The five capstones, at a glance

File	Tier	The sample project	What it proves
<code>undergraduate-pursuit-capstone.md</code>	B.S.	From NOFO to Submission-Ready Pursuit — a small HUBZone firm runs a fictional federal services recompute through the full pipeline in 14 weeks	The graduate of the B.S. can read a solicitation, score a pursuit, build an ORBITAL, write a capture plan, and submit a compliant package — end to end

File	Tier	The sample project	What it proves
graduate-capture-capstone.md	M.S.	Capture Management Engagement, \$10–50M — a mid-size firm’s capture manager runs a large, multi-prime IDIQ competition, including teaming, price-to-win, discriminators, color teams, and final-proposal-revision strategy	The graduate of the M.S. can <i>lead</i> a pursuit: structure the business, defend the price, govern the production, and steer the final competitive move
doctoral-research-capstone.md	Ph.D.	Empirical Study of the Calibration Engine & the Golden Thread — a research design testing the discipline’s own machinery: is the win-probability estimate honest? Does the doctrine learn?	The graduate of the doctorate can turn a practitioner claim into a falsifiable research question, design the study, gather the data ethically, and publish the answer
portfolio-capstone.md	All tiers / lifetime	The Lifetime Portfolio Arc — one continuous 20+ year arc tying the three tiers together	The learner can see their own career as a single pursuit — a Dream that becomes a World — and can steward the discipline

The README you are reading now is the fifth piece: the **how to run a capstone program** guide and the **common assessment rubric** that every tier shares.

How to run a capstone program

A capstone is a program, not an assignment. It needs a calendar, a rulebook, a review mechanism, and a defense. The doctrine's own governance discipline supplies the shape: **criteria written before the work begins** (doctrine/04).

The five rules that make a capstone honest

1. **Real documents, safe exercises.** Every capstone is built on a **real, published, closed** solicitation — never a live one. Nothing a student produces may bear on an active competition. The sample projects in this folder use *fictional* agencies and requirements precisely so they can be worked in a classroom without any risk of touching the real market; a student's own capstone should use a real closed solicitation (SAM.gov, Grants.gov, or the SBIR/STTR portals) unless the instructor explicitly chooses the fictional sample as the teaching object.
2. **The rubric is public.** The 4-level scale below is published before the first class. Students see exactly how their work is judged before they do it — the doctrine's gate discipline applied to grading.
3. **The defense is real.** Every tier ends in a **defense before a skeptical panel** — instructor, practitioners, and (where possible) an outside reader. The defense is where the doctrine is tested under questioning, and where a student who knows the words but not the discipline is separated from one who can do the work.
4. **Criteria before the call.** Where a capstone asks for a bid/no-bid call, a price, or a research claim, the criteria for the call are written **in advance** and graded before the answer is known. A gate charter graded in Week 3, judged against the call made in Week 12, is the whole discipline in miniature.
5. **Integrity is a skill, not a rule.** No fabricated evidence, no invented awards, no misstated past performance. The post-mortem is graded on honest comparison, not on being right. A capstone that fabricates fails the assessment outright — the same standard the federal market enforces on everyone who touches it.

The stack-agnostic pledge

Every capstone at every tier obeys doctrine/08: **concepts are durable, tools are swappable**. Students may use paper, a whiteboard, a spreadsheet, or any word processor. A deliverable may *never* depend on a specific product or system, and no capstone requires knowledge of any particular software. When a sample project mentions a tool of any kind — a scoring sheet, a compliance tracker, a review log, a data store — it means *your organization's tooling*, whatever that is, and the concept is what is being tested. If a capstone deliverable would break because a vendor changed its product, the deliverable is testing the wrong thing.

Program formats that work

The same five capstones run at different cadences. Match the format to the learner:

Format	Cadence	Who	Notes
14-week semester	One milestone per week, defense in Week 14	Degree-track students (B.S., M.S.)	The default; the undergraduate and graduate samples are built on it
Compressed bootcamp	4–6 weeks, two milestones per week	Fellowship cohorts, continuing-education	Cuts the drafting weeks, keeps the gates and the defense
Thesis / dissertation arc	6–24 months	M.S. thesis, doctoral dissertation	The doctoral sample is built for this; the research design <i>is</i> the milestone sequence
Lifetime portfolio	Years, revisited at each tier	Practitioners and alumni	The portfolio sample; the arc is the capstone

Teams vs. individuals

- **Undergraduate:** teams of 3–4, each member owning at least one deliverable; the defense is individual.
- **Graduate:** individual or small team, with a named capture lead who *owns the whole pursuit*; the defense is individual.
- **Doctoral:** individual by definition — a dissertation is a solo contribution, defended publicly.
- **Portfolio:** individual, with a committee that includes a practitioner mentor.

The panel

Every defense panel should include at least one **practitioner** (someone who has sat in a source-selection or bid/no-bid room) and at least one **faculty or research reader** (someone who will test the *reasoning*, not the jargon). The panel's job is the doctrine's job: push until the student names the limits, then judge whether the work is *defensible*, not perfect.

The common assessment rubric

Every capstone in this folder is graded on the same four-level scale. The scale is deliberately coarse — four levels force a decision and discourage the “everyone gets a B” drift that destroys grading as a governance mechanism.

The four levels

Level	Meaning	What the work looks like
Not Ready	The deliverable does not yet do its job	Structure missing or wrong; facts wrong or fabricated; no visible doctrine; a reader cannot act on it
Developing	The shape is right, the execution is partial	Right structure with thin content; some parts strong, several parts weak; doctrine used but superficially; a reader can follow but not rely on it
Proficient	The deliverable is professional and defensible	Correct structure, complete content, doctrine applied correctly; evidence cited for claims; a reader could act on it
Distinguished	The work a senior professional or a researcher would be proud of	Everything at Proficient plus: insight beyond the assignment, a hard problem met honestly, a defense that survives a hostile panel, and a lesson the field or the firm could use

The shared dimensions

Every tier-specific rubric in this folder is built from these six dimensions, weighted differently by tier:

Dimension	What it tests	Not Ready	Proficient	Distinguished
Domain mastery	Can the learner read and operate the real documents and structures of the discipline	Misreads core documents; wrong structures	Reads the solicitation, the factors, the volumes correctly	Reads them with insight — sees the strategy inside the compliance
Doctrine use	Is the pipeline, the gate, the score, the ORBITAL, or the color sequence visibly doing the work	No visible doctrine	Doctrine applied correctly	Doctrine applied and <i>questioned</i> — the learner knows where it is right and where it is thin
Evidence	Are claims grounded in data, sources, or documented reasoning	Assertion with no source	Factors, prices, and calls argued with cited evidence	Evidence is weighted and triangulated; the learner names what the evidence cannot prove
Governance	Were the gates, criteria, owners, and reviews real — and written before the work	Gates absent or after-the-fact	Gates identified with criteria and owners	The gate charter is graded before the call; the learner is <i>bound</i> by it
Artifacts	Are the deliverables complete, coherent, and mutually consistent	Large omissions; volumes contradict each other	Complete and consistent; the four volumes tell one story	The artifacts tell one story and the learner can defend every seam

Dimension	What it tests	Not Ready	Proficient	Distinguished
Defense	Can the learner survive a skeptical panel and own the limits	Evasive or unsupported	Clear and credible under questioning	Survives a hostile board; names the limits before the panel finds them

How to apply the scale per milestone

Each sample capstone includes a **per-milestone rubric**: for every week's deliverable, the criteria and what each level earns. Two practical rules for graders:

1. **Grade the milestone, not the trajectory.** A student who nails Week 4 but misses Week 9 gets credit for Week 4 and a developing on Week 9 — the capstone is a pipeline, and a weak stage downstream is a real defect, not a rounding error.
2. **The composite is the capstone.** Weight the milestones to match the tier's priorities (the undergraduate sample weights the ORBITAL and the score; the graduate sample weights the capture plan and the price; the doctoral sample weights the design and the ethics plan), then convert the composite to the same four levels. The panel's defense grade is the final gate — it can move the composite by one level in either direction, exactly as a real readiness gate can stop a submission.

How the tiers connect: the one arc

The four sample capstones are not four separate assignments; they are **four renderings of one arc** — the discipline from first read to first research contribution:

Undergraduate	Graduate	Doctoral	Portfolio
read → score build the ORBITAL CONSUME	lead → govern win the competition APPLY	question → test prove the doctrine GENERATE	steward → teach see the whole arc STEWARD

The undergraduate capstone proves the learner can *do* a pursuit. The graduate capstone proves the learner can *lead* a pursuit as a business. The doctoral capstone proves the learner can *study* the pursuit as a science. The portfolio capstone holds all three together over a career — the same person, the same doctrine, twenty years deeper. The `portfolio-capstone.md` sample explains the arc; each tier's sample shows its own rung.

The one-sentence summary of this folder

A capstone is the doctrine's final gate: real work, written criteria, an honest score, a skeptical defense — and the four sample projects in this folder show what that looks like at the bachelor's, master's, doctoral, and lifetime-altitude rungs of the same ladder.