

Strategic Initiative Track — The Doctrine as an Operating System (7 sessions)

For students from Strategic Initiative programs. The core gave you the pipeline. This track zooms out: the doctrine as an operating system for an organization — how strategy, portfolio, and governance turn a pipeline into a deliberate machine.

Audience: Strategic Initiative / strategy majors. **Prerequisites:** shared-core.

Readings: doctrine/03, doctrine/04, doctrine/05, doctrine/07, doctrine/08, and the case study.

Session SI-1 — The doctrine is an operating system

Learning objectives. By the end of this session, students can: (1) explain how a doctrine differs from a set of tools; (2) argue why durable concepts should be encoded, taught, and governed; (3) map an organization's stated strategy onto the pipeline stages.

Session plan (60 min). - Open (5 min): "What does an organization need to decide *before* it can pursue well?" - Teach (20 min): doctrine/08. The two layers — concepts and tools. Why the doctrine outlives the software. Why the doctrine is the shared language between humans and machines. - Apply (20 min): take any organization's strategic plan (use Ravonics). Place its stated goals onto the pipeline stages: where is it sensing, qualifying, scoring, delivering? - Discuss (10 min): "What happens to an organization that mistakes its tools for its doctrine?" - Close (5 min): assignment.

Discussion prompts. 1. Defend: "a doctrine is only real if it can be taught to a machine." 2. What is lost when an organization has no written doctrine?

Homework / reading. Read doctrine/08. Deliverable: a one-page "doctrine statement" for an organization you care about — the durable concepts it must hold, written so a new hire could apply them.

Comprehension check. What are the two layers, and which is durable? Why is the doctrine the shared language between human and machine?

Session SI-2 — Portfolio strategy: the pipeline as a portfolio

Learning objectives. By the end of this session, students can: (1) explain the pipeline as a funnel with conversion at each gate; (2) reason about pipeline value and where leverage lives; (3) recommend a portfolio allocation across risk levels.

Session plan (75 min). - Open (5 min): “If pursuit capacity is scarce, what deserves it?” - Teach (25 min): doctrine/03 + doctrine/05. The funnel view, pipeline value, threshold discipline, the trade between safe and bold pursuits. - Apply (30 min): the portfolio exercise. Give teams a fixed “capacity budget” and a set of scored opportunities (the same ones from core Session 5). Allocate capacity across safe wins, mid bets, and long shots. Defend the portfolio. - Discuss (10 min): “Which portfolio would you defend to a board — and why?” - Close (5 min): assignment.

Discussion prompts. 1. Where is the biggest conversion loss in a typical pipeline? 2. Should an organization ever pursue a below-threshold opportunity? Under what conditions?

Homework / reading. Read doctrine/05. Deliverable: a one-page portfolio recommendation for Ravonics — capacity budget, allocation, and the reasoning at each gate.

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

Session SI-2b — RICE: the human decision layer for the pursuit portfolio

Learning objectives. By the end of this session, students can: (1) apply the RICE framework — Reach, Impact, Confidence, Effort — to score and rank a pursuit portfolio; (2) run a human-agent co-drive, where an analyst agent computes the scorecard and the leader owns the weights and the call; (3) set a re-scoring cadence that keeps a portfolio honest as intelligence arrives.

Session plan (75 min). - Open (5 min): “SI-2 allocated a capacity budget by feel. What is the decision method behind a good allocation?” - Teach (25 min): the RICE framework applied to GovCon. Reach — the footprint of the prize (contract ceiling, platform breadth, agencies touched). Impact — the strategic weight of winning (new capability, incumbency, market entry, credibility). Confidence — pWin: how likely we are to win, and why. Effort — the scarce capacity from SI-2 (production, capture, teaming, leadership attention). The score is $(\text{Reach} \times \text{Impact} \times \text{Confidence}) \div \text{Effort}$ — and the denominator is why ambition gets discounted. - Apply (30 min): the portfolio exercise, part two. Teams re-run their SI-2 allocation with a RICE scorecard. The instructor works the Ravonics exemplar (below); each team computes scores, sets a threshold, and defends the ranked portfolio. - Discuss (10 min): “The agent computes the scorecard. Who owns the weights — and why does it matter that they do?” - Close (5 min): assignment.

Worked portfolio example. Ravonics is a five-person firm — it cannot staff every pursuit, so it must rank. The leader’s analyst agent assembles the scorecard from the opportunity data, the past-performance record, and the cost

model; the leader has set the year's weights as policy — reach and impact balance, confidence must clear the pWin bar, effort is what capacity is scarce. The score for each pursuit is $(\text{Reach} \times \text{Impact} \times \text{Confidence}) \div \text{Effort}$:

Pursuit	Reach (1–10)	Impact (1–10)	Confidence (0–1)	Effort (capacity units)	RICE = (R × I × C) ÷ E
Existing-customer recompete (engineering set-aside)	6	6	0.8	20	$(6 \times 6 \times 0.8) \div 20 = 1.44$
STTR research award with INSTAR Lab	5	8	0.6	25	$(5 \times 8 \times 0.6) \div 25 = 0.96$
HUBZone set-aside quick win	3	5	0.75	12	$(3 \times 5 \times 0.75) \div 12 = 0.94$
Mega-prime subcontract (teaming play)	8	7	0.3	45	$(8 \times 7 \times 0.3) \div 45 = 0.37$
Hosted-intelligence IDIQ (prime)	9	9	0.25	60	$(9 \times 9 \times 0.25) \div 60 = 0.34$

The rank is not a surprise to anyone who has run the doctrine: the recompete protects the base, the STTR award builds the research lane INSTAR Lab makes winnable, and the HUBZone set-aside is cheap confidence. The two glamorous pursuits — the IDIQ prime and the mega-prime subcontract — score last, because a five-person firm pays the full effort while carrying low confidence into the room. That is the RICE insight: the denominator (effort) and the confidence discount are exactly what raw excitement ignores.

The leader sets the threshold as policy. Here: pursue 0.8; park and re-score between 0.5 and 0.8; no-bid below 0.5. That clears the recompete, the STTR award, and the set-aside for pursuit, and sends the two long shots to the shelf. A deliberate override is still allowed — a leader may push a below-threshold pursuit

when the win would build a capability that changes next year's portfolio — but an override must be named as judgment, never mistaken for arithmetic. That is the human-agent co-drive in one rule: **the agent computes the scorecard; the human owns the weights and the override. The weights are policy, not output.**

Confidence is the most volatile input. The STTR award's confidence jumps the day INSTAR Lab confirms the research letter; the set-aside's falls if a stronger HUBZone competitor signals entry. So the portfolio is re-scored on a cadence: at discovery triage, at every gate (bid/no-bid, threshold, readiness), whenever a material signal lands, and in full at the quarterly portfolio review. A pursuit's RICE number is a living figure, not a plaque.

Discussion prompts. 1. What does the denominator do to a glamorous pursuit with low win confidence? 2. When is a below-threshold override the right call, and how do you keep it honest?

Homework / reading. Read doctrine/05. Deliverable: a RICE scorecard for Ravonics' portfolio — the weights you set as policy, the score for each pursuit, your stated threshold, and the ranked call you would defend to a board.

Comprehension check. What are the four RICE inputs, and what does each mean for a GovCon pursuit? Why are the weights policy, not output? When is a pursuit re-scored, and which input is the most volatile?

Session SI-3 — Gate governance: who decides, and how

Learning objectives. By the end of this session, students can: (1) design a gate with all four parts; (2) distinguish owner-with-standing from owner-with-process; (3) build the ladder of governance for a full pursuit.

Session plan (60 min). - Open (5 min): “Why does excitement make gates necessary?” - Teach (20 min): doctrine/04. The four parts of a gate, the two ownership modes, the ladder of commitment, and the ethics layer. - Apply (20 min): design the gate structure for a fictional pursuit. Teams write the criteria for the bid/no-bid gate, the threshold gate, and the readiness gate for a real opportunity. - Discuss (10 min): “Which gate in your design would fail first under pressure? How do you make it boring?” - Close (5 min): assignment.

Discussion prompts. 1. What makes a gate a rubber stamp, and how do you prevent it? 2. Where does public-side governance (the law) meet private-side governance (your gates)?

Homework / reading. Read doctrine/04. Deliverable: a one-page gate charter for one pursuit — the four parts written out for each gate the pursuit will cross.

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

Session SI-4 — The ecosystem view: partners, teaming, and the market

Learning objectives. By the end of this session, students can: (1) explain why no pursuit is an island; (2) reason about teaming and partnerships as strategy; (3) read the competitive landscape of a market using NAICS and past awards.

Session plan (75 min). - Open (5 min): “Who wins federal work — companies, or coalitions?” - Teach (20 min): doctrine/01 + doctrine/07. The market structure, the small-business universe as a design, teaming as strategy (the STTR research-partner model as an example). - Apply (30 min): for the Ravonics case, map its market. Use NAICS codes and published past awards to sketch who competes, who teams, and where a HUBZone-certified, research-partnered firm has an edge. - Discuss (10 min): “When is a competitor also a teammate?” - Close (5 min): assignment.

Discussion prompts. 1. Why does Congress require small businesses to partner with research institutions for STTR? 2. How does a small firm decide *who* to team with?

Homework / reading. Read doctrine/01 and the Ravonics case. Deliverable: a one-page market map for Ravonics — competitors, partners, certification advantages, and one teaming recommendation.

Comprehension check. Why is a firm’s certification posture a strategy, not paperwork? What makes the STTR model structurally different?

Session SI-5 — Organizational posture: DREAM, GROWTH, and SuperNova

Learning objectives. By the end of this session, students can: (1) explain the DREAM and GROWTH accents and where each lives in the doctrine; (2) explain the SuperNova arc from seed to operating entity; (3) argue which posture fits which organizational stage.

Session plan (60 min). - Open (5 min): “How does a seed idea become an operating entity?” - Teach (25 min): doctrine/03 + doctrine/07. The accents: DREAM for engagement, GROWTH for production, SuperNova for the full arc. Posture as a stage-dependent choice. - Apply (20 min): teams place an organization’s journey on the SuperNova arc. Where is it now? What is the next stage’s work and gate? - Discuss (10 min): “Which posture is most dangerous when applied at the wrong stage?” - Close (5 min): assignment.

Discussion prompts. 1. Why does a startup-stage entity need different governance than a mature one? 2. Where does the human (the DREAM relationship) matter most?

Homework / reading. Read doctrine/07. Deliverable: a one-page posture assessment for Ravonics — which stage of the arc it is in, and the three moves that would advance it.

Comprehension check. What is the difference between DREAM, GROWTH, and SuperNova as accents of one doctrine?

Session SI-6 — Leading a pursuit organization

Learning objectives. By the end of this session, students can: (1) assemble a pursuit organization from the doctrine's roles and gates; (2) define the metrics that tell a leader whether the pipeline is healthy; (3) present a doctrine-based strategy to a skeptical board.

Session plan (75 min). - Open (5 min): “What would you measure to know your pursuit machine is healthy?” - Teach (20 min): synthesis across doctrine/03–07. Roles, gates, portfolio metrics, the learning loop. - Apply (30 min): the board exercise. Each team plays the leadership of Ravonics presenting next year's pursuit strategy — pipeline posture, portfolio allocation, gate governance — to a skeptical board (the instructor + other teams). The board pushes back. - Discuss (15 min): “What made the strongest argument: numbers, structure, or story?” - Close (5 min): assignment and bridge to the capstone.

Discussion prompts. 1. What three metrics would you watch weekly? 2. How do you defend a “no” to a room that wants to chase?

Homework / reading. Read the full doctrine folder if you have not. Deliverable: a one-page “leader's dashboard” for a pursuit organization — the metrics and the gates that matter.

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