

A Day in the Life: From Solicitation to Bid/No-Bid

A complete, end-to-end walk of one realistic pursuit — read the document, build the compliance matrix, shape the ORBITAL, compute the score, and make the call. Every step shows the reasoning, not just the result.

What this is. This is the single most useful artifact in the practice layer: a worked example of the whole discipline on one pursuit, the way a real capture lead would actually do it. It is a **teaching object**. The solicitation below is a simplified composite built to look and behave like a real federal NOFO — it is not a live opportunity. Real documents are messier, longer, and harder to read than this one; that is exactly why this exemplar exists: to show you the *shape* of the work so a real document does not intimidate you.

Who. We are walking through the pursuit 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**.

What you will produce by the end. The five artifacts a professional capture team produces before committing to a pursuit: a bid/no-bid recommendation, a score, a compliance matrix, an ORBITAL, and the reasoning that holds all four together. These are the same deliverables as the capstone — the exemplar is a complete worked answer to that assignment.

The day, in one glance

MORNING	SENSE + QUALIFY	The 15-minute scan. Is this real, eligible, worth it? → YES
MIDDAY	SCORE	$p_{\text{Win}} \times \text{value}$, factor by factor, against a threshold → $0.77 \times \$250\text{k}$
AFTERNOON	COMPLIANCE + ORBITAL	The compliance matrix + the seven-axis ORBITAL
EVENING	THE CALL	BID. Written defense, in one paragraph. → BID

The rest of this page is that day, hour by hour.

The document that started it all

Every pursuit begins with a document. Here is the relevant heart of it — abridged, but structurally honest to how a real **STTR Phase I NOFO** reads. Keep this open; every step below quotes from it.

U.S. DEPARTMENT OF ENERGY, OFFICE OF SCIENCE Small Business Technology Transfer (STTR) Program — Phase I Release

A. Program Description. The DOE STTR Program supports cooperative research and development between small businesses and nonprofit research institutions. Phase I awards establish the *technical merit, feasibility, and commercial potential* of a proposed R&D effort. This Release makes funding available under the following topic:

Topic 2 — Physics-Informed Digital Twins for Energy Distribution Infrastructure. The Department seeks proof-of-concept research on physics-based digital twin models of electric power distribution systems that fuse real-time sensor data with physics models to predict equipment failure modes, optimize maintenance, and reduce outage risk. Successful Phase I efforts will demonstrate a prototype on a representative distribution feeder and articulate a credible Phase II path to a deployable product.

B. Award Information. Approximately \$55 million is available across all topics in this Release. For Topic 2, DOE anticipates 6–10 Phase I awards. **Each Phase I award is a fixed amount of \$250,000** for a **6-month** period of performance. Successful Phase I performers may apply for **Phase II** (up to \$1,100,000) in a later, separate competition.

C. Eligibility Information. 1. The applicant must be a **U.S. small business concern** (per SBA size standards) with an **active SAM.gov registration**, including a current **UEI** and **CAGE code**, at the time of application. 2. The applicant must have a **nonprofit research institution as a subawardee** (the STTR requirement). 3. The **small business must perform at least 40% of the total effort**; the **research institution must perform at least 30%**; the remainder may be performed by the small business or other partners. 4. The **Principal Investigator (PI) must have a primary employment with the small business concern**. 5. HUBZone, 8(a), and other certified small businesses are eligible and encouraged to apply.

D. Application and Submission Information. Applications are submitted electronically through the **DOE STTR/EE Submission Portal**. **Questions deadline: 2026-06-22, 5:00 PM**

ET. Submission deadline: 2026-07-15, 11:59 PM ET. Late submissions are rejected; the portal closes mechanically. Applications must include a **Technical Volume** (max **15 pages**, 12-point font, 1-inch margins), a **Budget and Budget Justification** using the provided template, a **Commercialization Plan** (max 3 pages), and completed **representations and certifications**.

E. Application Review Information. Applications will be evaluated on the following criteria and weights:

Criterion	Weight
Technical Merit and Innovation	40%
Impact and Significance	25%
Team and Qualifications	20%
Budget and Feasibility	15%

The review will assess whether the proposed work is technically sound and innovative; whether it addresses a significant problem; whether the team is qualified to perform it; and whether the budget and plan are realistic and aligned with the work.

F. Award Administration Information. The award will be subject to the STTR patent-rights clause (37 CFR 401), allowing the small business to retain title to inventions with reporting obligations to DOE. Grantees must submit an initial progress report, a final technical report, and a final financial report per the award terms.

G. Agency Contacts. Questions must be submitted in writing via the portal. DOE posts all questions and answers as an amendment to this NOFO.

Step 1 — Morning: sense and qualify (the 15-minute scan)

The NOFO lands in the morning. Nobody reads it front-to-back yet. The capture lead does a **15-minute scan** to answer the four questions that decide whether the rest of the day (and weeks) are worth spending:

1. **What is being funded?** A 6-month, \$250,000 proof-of-concept for a physics-based digital twin of energy distribution infrastructure.
2. **Are we eligible?** Small business — yes. Active SAM — yes. Research institution partner — yes, INSTAR Lab. STTR work-share (40% small business, 30% research institution) — satisfiable. PI primary employment at the small business — we have the person. HUBZone — yes, and the NOFO explicitly welcomes HUBZone firms.

3. **When is it due, and can we realistically be ready?** Submission 2026-07-15; today is 2026-06-01. Six weeks. Tight but feasible for a small team working one pursuit.
4. **Does it fit what we do?** This is the moment to check the topic against the four NAICS codes and the bench. Physics-informed modeling, sensor data fusion, prediction of failure modes — this is squarely **541715** (R&D in physical and engineering sciences) and **518210** (data processing and hosting), with a strong **541512** (systems design) flavor. The team that founded Ravonics left defense and government labs precisely to do this kind of modeling.

The qualify decision (the first gate). The pipeline’s first gate asks: *is this a real, eligible, fitting opportunity?* The answers are four yeses. The pursuit is qualified. Note what the gate did *not* do: it did not decide to win, or commit money, or promise a proposal. It only decided that this pursuit has earned a closer look. **A disciplined “no” here would have been a professional outcome too** — but on the facts, this is a clear “yes, look closer.”

One more qualification check that small firms learn the hard way: **the two-deadline trap**. The questions deadline (2026-06-22) and the submission deadline (2026-07-15) are different, and there is an invisible third deadline — the *internal* deadline (we set ours for 2026-07-08) that leaves a full week for review, fixes, and the mechanical act of submission. Put both on the calendar before anything else.

Step 2 — Midday: the score ($\text{pWin} \times \text{value}$, argued factor by factor)

Qualification says “worth a closer look.” The **score** turns “worth a closer look” into a number. Per doctrine/05, the score is:

$$\text{Score} = (\text{probability of win}) \times (\text{estimated value})$$

The value

The value is the easiest number in this pursuit, because the government published it: **\$250,000** Phase I. (This is a real advantage of SBIR/STTR work — the award size is known in advance.) There is also a Phase II at up to \$1,100,000, but the gate discipline says: **score the pursuit in front of you**. Phase II is a future, separate competition; we note it as upside in the ORBITAL, not as value in the score.

The probability of win (pWin)

pWin is not a gut feel. It is a **factor-weighted estimate that has to survive a room**. Ravonics uses the five classic inputs from doctrine/05, adapted to this

competition. The capture team argues each factor, assigns a score 0–1, weights it, and adds:

Factor	Weight	Score	The argument (why this score)
Topic fit / mission alignment	25%	0.90	The topic is the company's founding bet: physics modeling + AI + sensor fusion. The capability statement practically writes itself. Not 1.0, because topic fit alone never wins.
Team and research partner	20%	0.85	PhD physicists, AI specialists, Azure/AWS architects — plus INSTAR Lab, which satisfies the STTR research-institution requirement and adds real physics credibility. The partnership letter is in hand.

Factor	Weight	Score	The argument (why this score)
Past performance and capability	20%	0.55	Honest weak spot. A strong bench but a <i>thin federal award record</i> — no prior Phase I win. Some relevant work for government-affiliated labs. This drags the score down and must be addressed in the proposal.
Technical approach (our own draft)	20%	0.80	A defensible concept with a credible prototype path. The risk is integration scope, not the physics.
Capacity and capture readiness	15%	0.70	Small firm; this is the one pursuit the team can afford to staff properly. A proposal plan and a draft outline already exist.

Weighted pWin:

$$\begin{aligned}
& 0.25 \times 0.90 + 0.20 \times 0.85 + 0.20 \times 0.55 + 0.20 \times 0.80 + 0.15 \times 0.70 \\
&= 0.225 + 0.170 + 0.110 + 0.160 + 0.105 \\
&= 0.77
\end{aligned}$$

pWin 0.77.

The composite and the threshold

$$\text{Score} = 0.77 \times \$250,000 = \$192,500$$

Now the gate. Ravonics’s threshold rule, set in advance (this is what makes it a gate and not a mood): **pursue only if pWin 0.42 and the composite \$105,000** ($0.42 \times \text{value}$) for a pursuit of this size. The 0.42 is the same line the wider ecosystem uses for its Dream Score — a deliberate, defensible bar, not an accident.

- pWin $0.77 - 0.42 \rightarrow$ **pass**
- Composite \$192,500 - \$105,000 $\rightarrow$ **pass**

The pursuit clears the threshold gate and advances to *pursue*.

Why argue the score out loud?

Because scores change when you argue them. Run the sensitivity: if past performance were scored 0.30 instead of 0.55, pWin drops to 0.72 — still above the line. If topic fit collapsed to 0.40 (a misalignment), pWin drops to ~0.645 — still above the line. The call is **robust**, which is exactly what you want from a score before you spend real money. The only scenario that flips this pursuit is a much smaller value: at a hypothetical \$100,000 Phase I, the composite would be \$77,000 — below the \$105,000 line, and the honest call would become **BID WITH CONDITIONS** or **NO BID** depending on the portfolio. Knowing *that* in advance is the point of scoring.

Step 3 — Afternoon: the compliance pass (the compliance matrix)

The score says the pursuit is worth competing for. Now the work gets mechanical and unforgiving: **the compliance reading**. Missing a required page, a signature, or a certification can kill a technically excellent proposal instantly — and it is the easiest failure for an evaluator to find. So the team builds the **compliance matrix**: every requirement (“shall”) in the document, where it lives, and where the response will satisfy it.

Here is the matrix for this NOFO. Notice that this is *our* tool, not the agency’s — it is the backbone of the later review gates (the pink team starts here).

#	Requirement (the “shall”)	Source	Where we will satisfy it	Owner
1	Active SAM.gov registration with current UEI and CAGE at time of submission	§C.1, p.4	Confirmed active before submission; UEI/CAGE cited on the cover page	Capture lead
2	Nonprofit research institution as subawardee	§C.2	INSTAR Lab letter of commitment; subcontract in the budget	Program manager
3	Small business performs 40% of total effort	§C.3	Effort split shown in the budget; 40% Ravonics floor stated	Program manager
4	Research institution performs 30% of total effort	§C.3	INSTAR Lab budgeted at 30% of effort	Program manager
5	PI has primary employment with the small business	§C.4	PI bio + employer statement in Technical Volume	Capture lead
6	Technical Volume 15 pages, 12-pt font, 1-in margins	§D	Compliance-checked template; page count on the final PDF	Proposal writer
7	Budget and Budget Justification on the provided template	§D	Budget volume; justification aligned to the work plan	Program manager

#	Requirement (the “shall”)	Source	Where we will satisfy it	Owner
8	Commercialization Plan 3 pages	§D	Commercialization volume	Capture lead
9	Representations and certifications completed	§D	Completed in the portal before submission	Compliance officer
10	Submission via the portal by 2026-07-15, 11:59 PM ET	§D	Internal deadline 2026-07-08; upload verified on 07-09	Capture lead
11	Questions submitted in writing via the portal before 2026-06-22	§G	Two clarifying questions drafted and submitted 06-15	Capture lead

The strategy pass. The compliance reading answers “*what is required?*” The **strategy reading** answers “*what does the evaluator actually care about?*” Read the criteria’s own language: Technical Merit 40%, Impact 25%, Team 20%, Budget 15%. The proposal’s structure must **mirror the weights** — the best pages go to the technical approach, then to significance, then team. And read the adjectives: the NOFO asks for proof of *technical soundness and innovation*, a *significant problem*, and a *realistic, aligned* budget. Two strategy conclusions jump out:

1. The **past-performance weakness** (pWin factor 0.55) is the biggest threat to the Technical Merit and Team criteria — the proposal must lead with the bench and the physics credibility (and INSTAR Lab) to compensate.
2. The **budget-and-work alignment** is an explicit criterion — the single easiest way to lose points is an inspiring narrative and a disconnected budget. They must be one story.

Step 4 — Afternoon: the ORBITAL (the seven-axis skeleton)

Qualified, scored, and compliance-mapped — the pursuit is now real. Time to give it a skeleton. Per doctrine/06, the ORBITAL is the single reference the whole pursuit points at: an objective, a team, a budget, indicators, a route, activities, and the conditions that must hold. Here is Ravonics’s ORBITAL for this pursuit.

Axis	Letter	What it holds for this pursuit
Objective	O	Deliver a Phase I proof-of-concept physics-informed digital twin of a representative distribution feeder that predicts transformer failure modes from fused sensor + physics data. Success: a demonstrable prototype and a credible Phase II thesis. Tie to the agency’s stated outcome: reduce outage risk.
Resources	R	Ravonics: 2 PhD physicists, 1 ML engineer, 1 Azure/AWS architect, 1 program manager (part-time). INSTAR Lab: principal investigator + 2 researchers (the required research-institution share). External: no hires needed.

Axis	Letter	What it holds for this pursuit
Budget	B	Award: \$250,000 (fixed). Cost structure: Ravonics labor \$150k (60%), INSTAR Lab subcontract \$75k (30%), equipment/materials \$15k, travel + indirect \$10k. Pursuit cost to compete: \$8,000 (capture + proposal, mostly already sunk in people time).
Indicators	I	Weekly: milestone % complete vs plan. Monthly: internal review-readiness score. Gate-specific: compliance matrix 100% by 07-05; pWin factors re-argued at 07-01 internal gate. Success indicator for the <i>pursuit</i> : a compliant, reviewed submission by 07-08.
Transport	T	The vehicle is the STTR Phase I award (fixed \$250k). Delivery path: prototype demonstrated at a DOE stakeholder demo; reporting through the DOE STTR portal. Channel to customer: the program officer + the review panel; post-award, the contracting officer.

Axis	Letter	What it holds for this pursuit
Activities	A	A 26-week plan, in GROWTH language: Gauge (weeks 1–2: finalize requirements, questions), Research (weeks 3–4: data + literature), Organize (weeks 5–6: team, evidence, prototype plan), Write (weeks 7–10: technical/commercialization/budget volumes), Target (week 11: internal review gates), Health (weeks 12–26: award and deliver, or post-mortem).
Logistics	L	What must be true for the work to happen: SAM active; INSTAR Lab subaward and IP agreement executed (STTR patent clause 37 CFR 401 — Ravonics retains title, INSTAR research rights defined); no clearances required (unclassified); all facilities in Ravenswood; portal credentials tested before deadline day.

Two things to notice. First, **the ORBITAL makes the score honest**: the value (\$250k) and the cost structure (\$235k to deliver, \$8k to pursue) live in the Budget axis, so nobody can claim the pursuit is valuable while refusing to say what it costs. Second, **the ORBITAL is executable**: the Activities axis is a schedule, and the Logistics axis names the one thing most likely to kill a small firm’s pursuit (the INSTAR subaward and IP agreement) before it is executed.

Step 5 — Evening: the call

The score is argued, the matrix is built, the ORBITAL is filled. Now the actual decision — the gate that turns all the work into a commitment.

The gate, written in advance: BID if pWin 0.42 and composite \$105,000 and no disqualifying compliance gaps. Otherwise BID WITH CONDITIONS or NO BID.

The evidence assembled: pWin 0.77; composite \$192,500; all eleven compliance items satisfiable; budget within award; the one condition is the INSTAR subaward, which is executable in the first two weeks.

The decision: BID.

The one-paragraph defense:

Ravonics should bid on Topic 2. The pursuit is squarely in our core capability (physics modeling + AI + sensor fusion), we are eligible with our HUBZone status and our INSTAR Lab partnership, and the award size is known (\$250,000 Phase I). Our honest score is $\text{pWin } 0.77 \times \$250,000 = \$192,500$, well above our 0.42 / \$105,000 threshold, and the call is robust to reasonable swings in the weakest factor (past performance). The largest risk — a thin federal award record — is addressable in the proposal by leading with our bench and the research partnership. We can staff it as our one active pursuit, comply fully (all eleven matrix items), and deliver on schedule. BID, with the condition that the INSTAR subaward and IP agreement are executed by the end of week two.

Step 6 — After the call: the pipeline keeps moving

A bid/no-bid is not the end; it is the moment the pipeline shifts into production. From here, the work is:

- **Design** (weeks 1–8): the technical, commercialization, and budget volumes — each one answering the criterion it maps to, in the order the criteria are weighted.
- **Audit** (week 11): the color-team sequence. Pink reviews compliance (the matrix, item by item); red reviews persuasiveness against the criteria; gold reviews the price and the executive story. An audit failure here is a gift — it caught the problem while there was still time.
- **Submit** (by 07-08, a full week before the portal closes): the reviewed, compliant package goes in early, and the submission is verified, not assumed.
- **Deliver and learn** (Phase I award, or post-mortem): whether Ravonics wins or loses, the outcome gets studied. If it wins, delivery becomes the past performance that raises the next pWin. If it loses, the actual award

(published on SBIR.gov) shows who won, and the factor scores get updated with what that teaches.

That last move — reading the published award after the competition closes and updating the factor scores — is the **learning loop** (doctrine/03, stage 9), and it is the moment the day in this exemplar becomes a career skill. Do not skip it.

The lessons, in plain language

1. **The scan comes first.** Six weeks of work are saved or spent by a disciplined 15-minute read of four questions: what, who, when, fit.
2. **Qualify before you score.** The bid/no-bid gate and the threshold gate are different gates, and both are cheaper than a proposal.
3. **Scores are argued, not felt.** A pWin that cannot be defended factor by factor is a guess; a score that survives a sensitivity check is a decision.
4. **Compliance is a matrix, not a memory.** Every “shall” gets a row, a source, a satisfaction point, and an owner — before anyone writes prose.
5. **The ORBITAL is the spine.** It holds the budget, the plan, and the conditions in one place so the pursuit survives the week-to-week chaos.
6. **The call is boring by design.** Criteria written in advance, evidence assembled, owner accountable — the dramatic meeting dissolves into a check of boxes (doctrine/04).
7. **The loop is the point.** The day does not end at submission; it ends at the post-mortem, because the next pursuit starts with what this one teaches.

Try it yourself

Take any published, closed solicitation from SAM.gov or Grants.gov and run it through the same five steps: the 15-minute scan, the factor-by-factor score, the compliance matrix, the ORBITAL, and the written call. Then compare your call to what actually happened. That is the capstone, and now you have a worked answer to model yours on. The practice bank has graded drills at every step along the way.