

DPA26BZ05-DV019
Semantically-Aware ISR
FAQs

Question 1 – Direct-to-Phase-II Phase I-Equivalent Eligibility - The topic states that proposers shall provide documentation that Phase I-equivalent feasibility has been achieved outside the SBIR program and identifies five specific feasibility evidence elements.

Must an offeror have previously demonstrated the complete mission-aware semantic communications capability described in the topic, including each of the five listed evidence elements, or may Phase I equivalency be established through closely related prior R&D demonstrating the underlying perception, embedded edge processing, temporal tracking/reasoning, mission-relevant information extraction, and UAS integration capabilities, supplemented by analysis demonstrating the feasibility of the proposed semantic communications approach?

- A. Proposers must either be able to directly show Phase I feasibility related to all five areas or offer an explanation how their previous experience is applicable to the Phase 1 feasibility requirements.

Question 2 – Prior ROI/Semantic Transmission and Temporal/Contextual Reasoning -

The Phase I-equivalent criteria require a prior prototype that ingests EO video and produces ROI or semantic transmissions, including at least one task in which ROI identification depends on behavior change, persistence, or context rather than per-frame object class. Would a prior embedded ISR prototype that performs onboard perception, multi-frame target tracking, track persistence/reacquisition, and generation of mission-relevant target/track information satisfy this criterion if adaptive semantic packetization and communications-bandwidth allocation were not explicit objectives of the prior effort? Additionally, can multi-frame track persistence, reacquisition, or identity maintenance based on temporal observations satisfy the "behavior change, persistence, or context" requirement, or is higher-level behavioral/scene-context reasoning required?

- A. No, as described, this is ATR and not semantic communications. Proposers must demonstrate “a prior prototype that ingests EO video and provides ROI or semantic transmissions.”

Question 3 – Data Reduction and Baseline Measurements - The topic requires preliminary measurements showing a credible path to ≥ 90 percent reduction in transmitted ISR video data while preserving mission-relevant information, as well as preliminary comparison against relevant baselines. Must ≥ 90 percent transmitted-data reduction have been experimentally demonstrated prior to proposal submission, or may measured results from an existing edge-ISR prototype, combined with quantitative analysis of ROI/semantic transmission, establish the required credible path to the threshold? Similarly, must prior work have experimentally compared the prototype against all listed baseline categories (H.264/H.265/AV1, object detection/tracking, learned video compression, and lightweight vision-language perception), or are comparisons against the baselines relevant to the proposed technical approach sufficient?

- A. As the SBIR topic specifics, results should be measured and demonstrated, not just measured. Additionally, comparison against equivalent categories is acceptable.

Question 4 – Embedded Hardware / TRL 4 Requirement - The topic states that the projected TRL at the start of Phase II shall be 4 or above for the perception, reasoning, and semantic encoding subsystems and requires measured power and latency on an embedded inference target representative of small UAS payload compute. May prior measured embedded performance of perception/tracking/reasoning functions on representative sUAS hardware establish the required edge-deployment maturity, with semantic encoding and packetization subsequently integrated during Phase II, or must the semantic encoding subsystem itself have been implemented and demonstrated at TRL 4 or above prior to proposal submission?

- A. The semantic encoding subsystem must have been implemented and demonstrated at TRL 4 or above prior to proposal submission.

Question 5 – Supplemental Pre-Proposal Feasibility Work - May an offeror combine results from prior Government-sponsored R&D, including a prior SBIR Phase II addressing a different Government requirement, with company-funded feasibility testing performed prior to proposal submission to establish the Phase I-equivalent evidence required by this Direct-to-Phase-II topic? If so, may such supplemental testing be used to demonstrate specific elements such as semantic/ROI transmission, transmitted-data reduction, temporal/context-dependent ROI selection, baseline comparisons, and embedded power/latency when the underlying perception, tracking/reasoning, edge-compute, and

UAS-integration technologies were developed and demonstrated under the prior R&D effort?

- A. Yes, as long as it can be directly mapped to the Phase 1 feasibility requirements. Yes, subsequent testing can be used to demonstrate the other elements such as semantic/ROI transmission.

Question 6: How should offerors represent operator-defined mission intent for evaluation: natural-language mission statements, structured mission objectives, symbolic predicates, regions/times of interest, or a combination of these?

- A. It is up to the offeror to select the representation that works for their system and is most applicable to operator systems.

Question 7: What constitutes “preserved mission utility” for the 90 percent threshold and 95-99 percent objective data-reduction metrics, and which primary metrics will determine success (for example, event recall/precision, track continuity, time-to-cue, operator score, or downstream task performance)?

- A. Underneath Table I of the solicitation, the following statement applies:
“Mission utility should be evaluated using a combination of operator scoring, ROI/event recall and precision, track continuity, time-to-cue, semantic reconstruction quality, and downstream task performance relative to the original sensor stream.”

Question 8: What degree of object agnosticism is required: should the system discover previously unseen entities/events without predefined classes, or is it acceptable to use pretrained foundation models as long as region/event selection is not limited to a fixed mission-specific label set?

- A. The system is not expected to detect unseen objects but is expected to detect and identify things pretrained on.

Question 9: What multimodal sensor combinations should be considered mandatory for the principal demonstration (electro-optical, infrared, event camera, low-light, telemetry), and must all modalities be processed concurrently on the same low-size-weight-and-power payload?

- A. The only required modality is EO. Offerors should propose additional modalities and “by the end of Phase II, demonstrate EO video plus at least on additional modality or metadata source”. All modalities should be processed on the same payload.

Question 10: What degraded/disrupted/intermittent/limited communications profiles should Phase I and Phase II use, including expected bandwidth, latency, packet loss, outage duration, and emission-control constraints?

- A. While no specific DDIL profiles are outlined in the solicitation, offerors should discuss what profiles they plan on using and should show how they are indicative of real-world, operational values. At a high level, assume bandwidth lower than 10Kbps.

Question 11: Does the approximately 2-watt incremental compute objective apply only to the semantic-selection/encoding function, or to the entire onboard perception, reasoning, packetization, and communications-processing chain?

- A. Refer to the metrics, specifically Edge power, for more information.
 - a. Threshold: $\leq 5W$ on COTS embedded prototype
 - b. Objective: $\leq 2W$ incremental power module excluding camera and radio

Question 12: What traceability artifacts are expected for each transmitted semantic packet so that an operator can relate selected regions/events back to original sensor evidence and understand why the information was retained or suppressed?

- A. It is up to the offeror to specify which artifacts they can provide. The offeror must meet the requirements of the “**Traceability and operator trust**” under the Phase II Performance metrics section.

Question 13: What Group 1 and Group 2 unmanned aircraft system compute interfaces and payload envelopes should be treated as representative, and are there preferred processors or evaluation boards for benchmarking?

- A. Offerors are encouraged to use the systems listed in the solicitation as reference: “Small unmanned aircraft systems (sUAS) embody this challenge in its most acute form. Representative Group 1 platforms such as RQ-11B Raven and RQ-20 Puma-

class systems carry EO and IR video for dismounted operators with single-digit-watt payload compute budgets. Representative Group 2 platforms such as ScanEagle-class systems provide broader reconnaissance...”

Question 14: What baseline video codecs, region-of-interest methods, event-triggered approaches, or semantic-compression methods will the Government use as comparison baselines?

- A. Baseline video codecs include H.264, H.265, and AVI. Other baseline methods will be discussed with the Government once the effort begins.

Question 15: For the optional multi-unmanned-aircraft-system extension, should semantic coordination include cooperative task allocation and cross-platform information prioritization, or is the emphasis primarily on communications-efficient fusion and de-duplication?

- A. The option should “Extend the system to semantic coordination across two or more UAS feeds. Demonstrate cross-platform track and ROI consistency, semantic de-duplication, and prioritized transmission across overlapping or complementary mission regions.

Question 16: For the Phase I compression feasibility prototype, must the SWaP conditions described be met, or is a demonstrated path to meeting the SWaP provision sufficient?

- A. A path to meeting the SWaP provision is sufficient.

Question 17: Would documented edge deployments of perception tasks to platforms such as Team Awareness Kit (TAK) phone platforms be responsive, or are only deployments to UAS platforms of interest.

- A. This SBIR is seeking integration, test, and demonstration on “low-SWaP ISR platforms.”

Question 18: Can non-SWaP constrained platforms be presumed to be available for construction of mission cards targeting low SWaP platforms, or must the translation of vague natural language objectives happen on the low SWaP platforms?

- A. The focus is on Group 1 or Group 2 UAS and by the end of Phase II, “performers shall demonstrate integration with both representative classes of Government-approved surrogates.”

Question 19: For the Month 12+ sUAS integration demonstrations, will the government supply the platforms and range operations or should they be proposed?

- A. This will be arranged later in the program based on performance results and interested transition/integration partners.

Question 20: Which agency, office, and/or major program is the intended transition partner of the capabilities developed on this program?

- A. The performer and Government team will work together during the SBIR effort to identify the appropriate transition partner(s).

Question 21: Is it desirable or expected to have a University as part of a performing team? Is it possible to structure some of this effort as fundamental research? Will the publication of at least some of the results on this effort be possible?

- A. A university can partner on this effort. It is possible to structure part of the research as fundamental but requires contracting officer and program manager approvals. Publication of results will require approval from the DARPA Public Affairs Office (there is a process that is contractually required for all but fundamental research).

Question 22: Are you more interested in adapting existing approaches vs. pursuing new innovations?

- A. DARPA SBIR topics are looking for innovative approaches vs. incremental improvements to the state-of-the-art.

Question 23: Have you been working with others on projects related to this topic? Are there related SBIR topics, existing DARPA efforts, or other programs that would affect the scope of this topic?

- A. This is a new SBIR topic but it is related to the DARPA FLUID program.

Question 24: How would you compare and contrast the concept of “saliency” (foreground vs. background) with “semantic-awareness” within the context of a mission?

- A. Saliency answers the question of “What stands out?” while semantic-awareness within the context of a mission answers the question “What is the operational meaning of this image?”.

Question 25: How many potential mission intents do you envision for Phases I or II? Could mission intents be organized into distinct enumerated concepts or are they more like generalized categories? Are mission intents intended to be sensor or domain-specific?

- A. Distinct enumerated concepts initially with the potential of generalizing based on performance results.

Question 26: Are the listed Jetson, Hailo, Coral, FPGA, and Cortex-M targets illustrative, or will the Government favor a particular deployable supply chain? May perception and symbolic reasoning be processed on multiple or separate low-power devices?

- A. These are representative embedded targets only. It’s possible to deploy this on any of the types listed but it must meet the power requirements outlined in the solicitation

Question 27: Who defines the mission-utility function, and how are operator scoring, event recall, track continuity, time-to-cue, uncertainty calibration, and reconstruction quality combined? Will performance be averaged, scenario-weighted, or constrained by worst-case missed events?

- A. Specifics will be discussed once the effort begins. From the SBIR: “Mission utility should be evaluated using a combination of operator scoring, ROI/event recall and precision, track continuity, time-to-cue, semantic reconstruction quality, and downstream task performance relative to the original sensor stream.”

Question 28: Will the Government provide a mission-card or semantic-packet schema, or should performers define one? Should it interoperate with existing formats and tactical ISR standards?

- A. Performers should define this and it should operate with existing standards.

Question 29: Should semantic packet priority and evidence metadata remain visible to downstream networking middleware so it can selectively drop or aggregate packets, or should network packets be opaque and protected end-to-end?

- A. Packets need to be routable and processed by any downstream packet processing pipeline. Downstream routers do not need to understand the semantics.

Question 30: The topic requires "threshold = 90 percent (10x) reduction versus full-frame mission-video baselines," with a Table 1 objective of "≥ 95-99% reduction across mission cards." Is the intended baseline the raw uncompressed sensor stream, or an operationally representative encoded baseline such as H.264, H.265 or AV1 at a stated bitrate? If encoded, will DARPA specify the codec, resolution and bitrate for evaluation? The distinction changes reported reduction by more than an order of magnitude and is central to how we design and report against Table 1.

- A. The intended baseline is the already compressed video stream. Representative codecs are listed in the SBIR topic and DARPA will help specify details once the SBIR starts.

Question 31: The topic states that "semantic packets shall combine high-fidelity regions of interest, tracklets, scene-change descriptors, temporal evidence, low-rate background context, confidence, uncertainty, and traceable explanations."

For the data-reduction metric, should transmitted bytes be accounted inclusive of ROI imagery, periodic background context, transport and fragmentation overhead, forward error correction, and retransmissions? We intend to report fully inclusive totals and wish to confirm that is the intended basis rather than, for example, semantic metadata alone.

- A. The intent is a fully inclusive total.

Question 32: The low-SWaP embedded execution bullet states "Threshold = 5 W incremental power during interim testing; final objective = 2 W average incremental power excluding camera and radio."

Table 1 states "≤ 5 W on COTS embedded prototype," and the Month 9 milestone states "Threshold = 5 W on COTS embedded hardware."

For the 5 W threshold, which basis governs scoring: incremental draw attributable to the capability, or total board input power? Relatedly, is measurement on a COTS development kit acceptable, or is a production payload module expected? Development kits carry carrier-board, cooling and storage overhead that a payload module does not, and the two readings can differ by several watts on the same workload.

- A. The intent is that the solution can run on a board that outputs 5W total power. Measurement on a COTS development kit is acceptable.

Question 33: The Month 18 milestone requires "≥ 90 percent data reduction at iso-utility, with 20 to 30 dB stretch in mission-suitable scenes." Could DARPA clarify the intended definition and measurement basis for this stretch figure — for example whether it refers to a link-budget or signal-to-noise margin, a ratio on transmitted data volume, or a reduction in the load offered to the communications subsystem?

A. It refers to the reduction in the load offered to the communications subsystem.

Question 34: The topic states that mission utility "should be evaluated using a combination of operator scoring, ROI/event recall and precision, track continuity, time-to-cue, semantic reconstruction quality, and downstream task performance relative to the original sensor stream."

At Phase II evaluation, will operators be Government-provided, or is the offeror expected to recruit and fund a rater panel? Is an existing scoring rubric or evaluation protocol available for proposers to reference, so that offeror-run pilots are consistent with Government evaluation?

A. The Government will work with the performers to identify operators once the SBIR begins.

Question 35: The topic states that "Government-furnished platforms are not required for proposal submission" and that performers "may use representative surrogate airframes, software-in-the-loop, hardware-in-the-loop testbeds, recorded UAS video, captive-carry tests, and emulated degraded links."

- For the Group 1 and Group 2 integration milestones, will the Government furnish interface control documents, SDK or flight-controller API access for the named platform classes, or is the offeror expected to procure and integrate independently?
 - The offerer is expected to procure and integrate independently
- What constitutes a "Government-approved surrogate" for the Month 18 capstone, and at what point is surrogate approval sought — at the System Requirements Review, or later in the period of performance?
 - The Government will discuss approval of surrogates during the course of the program. Initial ideas of possibilities for surrogates should be discussed at the System Requirements Review.
- Are surrogate-based demonstrations evaluated equivalently to demonstrations on the named platform classes, or is a specific class expected for at least one of the two required integrations?

- Surrogate-based demonstrates are equivalent.

Question 36: The topic lists "long-wave infrared, low-light video, event-camera streams, platform telemetry, gimbal or inertial state, acoustic cueing, or RF or link-state metadata" as candidate second modalities.

Are all of these equally acceptable for the EO-plus-one-modality requirement, or are some preferred? Specifically, can a metadata-only stream such as platform telemetry, gimbal state or link-state metadata satisfy the requirement, or is an actively sensed second imaging modality expected?

- A. None are specifically preferred. From the SBIR topic: "Extend reasoning beyond single-modality EO video. By end of Phase II, demonstrate EO video plus at least one additional modality or metadata source: long-wave infrared, low-light video, event-camera streams, platform telemetry, gimbal or inertial state, acoustic cueing, or RF or link-state metadata."