DPA27BZ02-DV002
Autonomous Metacognitive Collection Orchestration for Multi-Domain Intelligence
Frequently Asked Questions (FAQs)

1. **May results from separate prior demonstrations collectively substantiate the three feasibility criteria, provided the package includes the required baseline comparisons and evidence of operation on at least one multidomain task with asynchronous inputs? Or should a single prior prototype have demonstrated all three criteria together?**
   Answer: Yes. Results from separate prior demonstrations can collectively substantiate the three DP2 feasibility criteria. The topic text states: *"DP2 feasibility shall be substantiated through prior work completed outside of the SBIR program... Proposers shall demonstrate scientific and technical merit and feasibility across the criteria below, each with quantitative side-by-side comparison against at least one mainstream baseline... on a shared task."* It does not mandate that a single prior prototype must have shown all three together, provided that the package includes a consolidated technical report summarizing all criteria with quantitative metrics, and provides video/data artifacts demonstrating operation on at least one multi-domain task with asynchronous inputs.

2. **May the benchmark task vary by criterion if the proposed method and its baseline are evaluated on the same task for each comparison?**
   Answer: Yes. The topic indicates that each of the three criteria must have a *"quantitative side-by-side comparison against at least one mainstream baseline... on a shared task"* for that specific criterion. There is no requirement that the exact same task must be used across all three distinct criteria, so long as the proposed method and its baseline are evaluated on the same shared task for each individual comparison.

3. **May the required prior demonstration use simulated data if it captures the specified differences in input timing and tests the required conditions for missing or delayed inputs?**
   Answer: Yes. Using simulated/synthetic data for prior feasibility work is acceptable, provided it realistically models the native tempo differences (two or more orders of magnitude) and tests the required conditions (e.g., at least 20 percent missing or delayed inputs).

4. **For results supported by an internal technical report and underlying test records, does DARPA expect independent validation of the reported comparisons?**
   Answer: No. The BAA does not explicitly expect independent external validation of the reported comparisons in the proposal stage. Proposers must substantiate feasibility using technical reports, prototype screenshots, test data, and a consolidated technical report outlining their own results. However, DARPA will perform a Policy and Security Review on selected proposals, and the evaluations are based on the documented evidence provided.

5. **May the three Direct to Phase II feasibility criteria be substantiated by different members of the proposing team, for example a deterministic constraint and audit component by the prime and the reasoning substrate by a subcontractor, each with prior work outside SBIR?**
   Answer: Yes. The feasibility criteria can be substantiated by different members of the proposing team (e.g., prime and subcontractor) as long as there is clear provenance and they are integrated into a single cohesive proposal. The BAA requires details on subcontractor pricing, SOW, and capabilities in the overall proposal (specifically in Volumes 2 and 5).

6. **For criterion 3, does DARPA accept as part of the "internal system state" a deterministic constraint layer, separate from the reasoning substrate, whose signed records reproduce each accept or decline decision and name the constraint applied?**
   Answer: Yes. A deterministic constraint layer that runs separately from the reasoning substrate, generating signed records of "accept" or "decline" decisions based on specific constraints, can be part of the "internal system state" to satisfy the auditability requirements of Criterion 3, provided that the system provides reproducible "why / why-not" traces directly from this state.

7. **Is a design that holds selected constraints (emissions-control posture, the hold-fire confidence threshold) at zero violations by construction, while the planner optimizes inside them, consistent with the topic's target of a violation rate below 5 percent?**
   Answer: Yes. Designing selected hard operational constraints (such as emissions-control or hold-fire thresholds) to have zero violations *by construction* is completely consistent with and exceeds the topic's target of a violation rate of *"below 5 percent under adversarial stress.*

8. **For criterion 1, is a demonstration on public data (ship-tracking messages at second-scale tempos and daily detections derived from overhead imagery) acceptable as two heterogeneous modalities whose native tempos differ by two or more orders of magnitude?**
   Answer: Yes. Fusing ship-tracking messages (e.g., cooperative AIS tracks updating at second-scale tempos) and daily detections derived from overhead imagery (updating at hours/days cadences) meets the requirement for *"two heterogeneous modalities whose native tempos differ by two or more orders of magnitude"* (second-scale vs. daily-scale is a difference of over 4 orders of magnitude.

9. **Multi-domain task: The DP2 feasibility package must show operation "on at least one multi-domain task with asynchronous inputs." Does a single maritime task that fuses two different sensing modalities with very different tempos (for example, cooperative AIS tracks and spaceborne SAR detections) satisfy this? Or must the feasibility task span two or more operational domains?**
   Answer: Yes. A single maritime task that fuses two different sensing modalities with very different tempos satisfies the requirement of operation on *"at least one multi-domain task with asynchronous inputs."* In military terminology (e.g., Joint All-Domain

Command and Control - JADC2), maritime surveillance is considered a multi-domain task since it crosses the space, air, and maritime domains.

10. **Archived-data replay: Is a prototype demonstration that replays archived, publicly released sensor data qualifying feasibility evidence? In such a replay the collection decisions and their bandwidth costs are simulated against what was actually recorded, rather than tasking live sensors.**
    Answer: Yes. A prototype demonstration that replays archived, publicly released sensor data (where collection decisions and bandwidth costs are simulated against recorded baselines) qualifies as valid feasibility evidence.

11. **Video artifacts: The topic asks for "video or data artifacts demonstrating operation," but Volume 2 may not embed video. Where should a video artifact be provided: as a file in volume 5, as a link, or only described in the white paper?**
    Answer: In Volume 5 or via links/descriptions. Since Volume 2 (Technical Volume) is uploaded as a static PDF and cannot embed active video elements (*"Do not include or embed active graphics such as videos..."*), you should describe the video in the white paper (Volume 2) and upload the video files or provide clear instructions/links in Volume 5 (Supporting Documents).

12. **CMMC timing: For a firm proposing ITAR work, must the CMMC Level 2 self-assessment be posted in SPRS by proposal submission, or by contract award?**
    Answer: Prior to selecting for award / Contract execution. The BAA states: *"Firms engaging in Controlled Unclassified Information (CUI), Export Controlled (EXPT), or International Traffic in Arms Regulations (ITAR) work for DARPA must have CMMC Level 2 self-assessment certification... please ensure the DARPA SBIR/STTR Pre-Award Checklist... is completed prior to being selected for award."* Having the score posted in SPRS is a prerequisite for award execution rather than the initial proposal submission date itself, though starting the process immediately is highly recommended.

13. **DP2 feasibility evidence: The topic requires DP2 feasibility substantiated through prior work completed outside the SBIR program, with quantitative side-by-side comparison against at least one mainstream baseline on a shared task for each of the three criteria (structured reasoning over asynchronous multi-source evidence; cost-aware forward-looking action selection; auditable decision outputs grounded in internal state). Will prototype demonstrations with data artifacts and a consolidated technical report - without peer-reviewed publications – be sufficient evidence if the quantitative baseline comparisons are included, or is a publication /pre-print record expected?**
    Answer: Yes. Peer-reviewed publications are *not* strictly mandatory. The topic text lists acceptable evidence as: *"peer-reviewed publications, pre-prints, patents, prototype demonstrations with video or data artifacts, technical reports, and test-and-measurement data."* A working prototype with video/data artifacts and a consolidated technical report containing quantitative comparisons is sufficient.

14. **Data, scenarios, and simulators: Will the Government provide, or designate, any data feeds, scenario sets, simulators, or an analytic / data-fabric environment for the DP2 feasibility demonstration or for Phase II evaluation, or should the offeror build and propose its own multi-modal simulators and scenario set for both? If Government-furnished material is anticipated, at what point (kickoff, or earlier) will it be identified?**
    Answer: No. The BAA does not promise or identify any pre-packaged GFM data feeds, scenario sets, or simulators before or at kickoff. The topic states: *"Baselines and the capstone scenario set will be finalized with the Program Manager at kickoff."* The offeror is expected to propose and build its own multi-modal simulators and scenario sets.

15. **Baselines and capstone finalization: The topic states baselines and the capstone scenario set will be finalized with the Program Manager at kickoff. For the proposal itself, should the offeror name its chosen baseline family (classical fusion, reinforcement / active learning, or agentic-LLM orchestration) and propose its own capstone scenarios from the three categories in the topic, subject to kickoff agreement?**
    Answer: Yes. The offeror must explicitly name its chosen baseline family (e.g., classical fusion, reinforcement/active learning, or agentic-LLM orchestration) and propose its own capstone scenarios in the proposal, which will then be finalized with the PM at kickoff.

16. **Traceability expectations: For the forward-looking "why / why-not" traces reproducible from internal system state, is there a preferred trace format, schema, or level of granularity DARPA expects (for example, evidence chain + constraints satisfied + dominated alternatives + forecast futures per recommendation), and will the SME coherence evaluation rubric (85 percent target) be shared at kickoff or should the offeror propose a doctrine-aligned rubric in the proposal?**
    Answer: There is no single mandated trace schema format, but it must contain the semantic elements listed (evidence chain, constraints, dominated alternatives, forecast futures). The 85% SME coherence rubric is finalized after kickoff, but the offeror is highly encouraged to propose its own doctrine-aligned rubric within the proposal.

17. **Hard operational constraints: The topic requires emissions, bandwidth, power, time-to-decision, and access posture to be treated as hard-binding semantic variables rather than reward-shaped scalars. Should the proposal demonstrate, in the feasibility package, at least one constraint enforced as a hard feasibility filter that can veto an otherwise higher-information-gain collection action - and is the below-5-percent constraint-violation target measured against that same hard-constraint definition under adversarial stress?**
    Answer: Yes. The proposal's feasibility package should demonstrate treating at least one constraint (e.g., emissions or bandwidth) as a hard semantic variable (veto filter). The Phase II constraint-violation target of below 5% is measured against these hard-constraint thresholds under adversarial stress.

18. **Quantitative targets and period structure: Please confirm the Phase II quantitative goals as we should plan against them - time-to-decisive-cue at least 35 percent faster**

than a reactive fusion baseline at fixed sensing budget; information-gain per unit energy at least 2x over a reactive fusion baseline; constraint-violation rate below 5 percent under adversarial stress; plan-refresh latency below 3 minutes; and SME trace-coherence at or above 85 percent. The BAA award-structure table lists $1,500,000 / 18 months with a $500,000 / 6-month option for this topic, while the topic text describes a 24-month Phase II with Month 21 and Month 24 milestones under "Option Milestones" - should we structure the base as Months 1-18 and the option as Months 19-24 accordingly?

Answer: The Phase II goals are:

a. **Time-to-decisive-cue:** ≥ 35% faster than a reactive fusion baseline.
b. **Information-gain per unit energy:** $\ge 2\times$ over a reactive fusion baseline.
c. **Operational-constraint violation rate:** < 5% under adversarial stress.
d. **Plan-refresh latency:** < 3 minutes.
e. **SME Trace-Coherence:** ≥ 85% against doctrine rubrics.

To align the BAA award structure table ($1,500,000 / 18 months base + $500,000 / 6 months option) with the topic's 24-month schedule, you must structure the Phase II Base period as Months 1–18 and the Phase II Option period as Months 19–24 (which includes the Option Milestones at Month 21 and Month 24).

19. **CMMC and classification: The topic projects CMMC Level 2 (self), and the BAA states firms engaging in CUI, export-controlled, or ITAR work for DARPA must have a CMMC Level 2 self-assessment. Is a completed Level 2 self-assessment (with score posted in SPRS) expected at proposal submission, at award, or by a later milestone?**

    Answer: At Award. The completed self-assessment and posting of the score in SPRS must be completed prior to contract award/execution. (additional information: You can submit if you don't have level 2 CMMC, but it is advisable to outline a plan for obtaining that certification level. You must have a self-assessment logged by time of award, and if you do not have a passing score, you need to provide a POAM (Plan of Action & Milestones) indicating how your firm plans to reach full compliance within 180 days of award.

20. **ITAR / foreign-national disclosure: The topic is ITAR/EAR-restricted and requires disclosure of any proposed use of foreign nationals, including country of origin, visa / work-permit type, and the SOW tasks they would perform. Our CTO is UK-based. If all ITAR-controlled technical work is performed in the United States by U.S. persons and any non-U.S.-person involvement is limited to clearly separated, unclassified commercial platform tooling with no access to topic technical data, what disclosure format does DARPA expect in the proposal, and are there tasks (for example, program management) that would still be treated as restricted?**

    Answer: Proposers must disclose the use of foreign nationals (such as your UK-based CTO) in Volume 7 (the Disclosures webform). If the CTO’s involvement is strictly limited to unclassified, non-export-controlled commercial platform tooling and they do not have access to any restricted technical data, you must clearly delineate this separation of duties in your SOW and Management Plan. Program Management tasks are generally considered restricted if they involve oversight of ITAR technical decisions or data.

21. **Deliverables and TRL entry: Phase II is described as advancing from TRL 4 at award to TRL 6 by the end of Phase II. What evidence of TRL 4 at award does DARPA expect in the DP2 feasibility package, and beyond the milestone briefings, quarterly reports, and Final Phase II Report listed in the topic, are there topic-specific deliverables (software drops, API specifications, manuals) expected at each milestone?**
    Answer:
    a. **TRL 4 Evidence:** Proven components validated in a laboratory environment (Criterion 1 & 2 demonstrated on synthetic or open datasets).
    b. **Specific Deliverables:** Beyond quarterly reports, DARPA expects software prototypes (deliverables at Month 12 and 18), API specifications, interface designs, user and administrator manuals (Month 24), and commercialization/transition plans.

22. **Can authenticated prior Phase-I-equivalent feasibility artifacts developed by a proposed subcontractor/team member be relied upon, with clear provenance, access/data rights and a defined continuing technical role, or must the prime itself have generated all prior demonstrations?**
    Answer: Yes. The prime contractor does not need to have generated all prior feasibility demonstrations. Feasibility artifacts from team members/subcontractors are fully acceptable if clear provenance, data rights, and their continuing role on the team are defined.

23. **Would a narrowly defined auditable constraint/reasoning subsystem, embedded in a teaming partner's proven two-modality orchestration demonstration, be considered within scope, provided the full team substantiates all mandated prior feasibility criteria? We understand a generic governance layer alone is insufficient.**
    Answer: Yes. Embedding a specialized auditable constraint/reasoning subsystem within a teaming partner's existing two-modality demonstration is appropriate and inside scope, provided the combined proposal clearly addresses all required feasibility criteria.

24. **The topic requires Phase I-equivalent feasibility from work outside the SBIR program, with quantitative comparison to a mainstream baseline. For a first-time small business, what documentation do you expect? Is a working software prototype with measured results on unclassified synthetic or open data sufficient, or are publications or patents also expected?**
    Answer: A working software prototype with measured quantitative results on unclassified synthetic/open data, documented within your Technical Volume (Volume 2) and accompanied by a technical report, is fully sufficient. Patents or peer-reviewed publications are not strictly required.

25. **What data types and sources, if any, will be furnished to a performer? Is a software-only, unclassified prototype demonstrated on synthetic or open multi-source data acceptable for the feasibility package and for Phase II evaluation?**

Answer: Yes. A software-only prototype demonstrated on synthetic, simulated, or open multi-source data is fully acceptable for both the feasibility package and Phase II evaluations. No hardware-in-the-loop is required, though the software must be designed for tactical-edge deployment constraints (such as Size, Weight, and Power - SWaP).

26. **The topic requires "why / why-not" traces reproducible from internal system state. How will explainability and audit trails be weighed in evaluation?**
Answer: Explainability and audit trails will be weighed significantly during evaluation. Under Criterion 3, generating forward-looking, reproducible traces directly from the internal state is a core technical requirement to ensure trustworthy human-machine teaming.

27. **What integration targets or end-user environments should a Phase II prototype be designed against?**
Answer: The Phase II prototype should be designed to integrate with standard military analytical/data-fabric environments (such as Joint All-Domain Command and Control - JADC2/CJADC2 frameworks) and tactical edge platforms.

28. **The topic requires an after-action audit interface and decision traces that are reproducible from internal system state. Can a subcontractor provide the audit interface and an independent check of the decision record as a separate component, or does DARPA expect the proposer to build both inside the reasoning engine itself?**
Answer: Yes. A subcontractor can build the after-action audit interface as a separate component, provided it integrates seamlessly via well-documented, open-schema APIs with the primary reasoning engine to retrieve the reproducible system states.

29. **Can feasibility evidence come from a non-ISR domain with the same structure, meaning asynchronous heterogeneous streams with cost-aware selection? Or must it use ISR or multi-INT data?**
Answer: Yes. Feasibility evidence can come from a non-ISR domain (e.g., commercial autonomous systems, industrial IoT, or commercial cyber networks) as long as it demonstrates the exact same structure of managing asynchronous heterogeneous data streams with cost-aware, constraint-bounded selection.

30. **Is simulated, synthetic, or public data, such as open AIS, acceptable for the feasibility package?**
Answer: Yes. Open AIS and other public datasets are highly acceptable for modeling asynchronous inputs and proving the feasibility package criteria.

31. **The topic is substrate-agnostic and encourages hybrids. Will a hybrid HDC/VSA and graph approach be judged on metrics for both families, such as retrieval and compositional capacity plus path-query accuracy, or must the proposer pick one primary family? For example, neuro-symbolic approaches, hybrid, often surface intelligence that is difficult for independent approaches.**
Answer: The topic is substrate-agnostic. If a proposer uses a hybrid approach (e.g., HDC/VSA and graphs), they should characterize the system using relevant metrics from

both families (e.g., retrieval/compositional capacity for HDC, and path-query accuracy for graphs).

32. **On "local refinement without offline retraining": would few-shot online updates to a system's memory, with no gradient-based retraining, meet this requirement? Would HDC (hyperdimensional computing) Prototype updates meet this requirement?**
    Answer: Yes. Few-shot online updates to memory models, or localized memory vector updates in HDC without gradient-based global retraining, fully satisfy the requirement for *"local refinement without offline retraining".*

33. **Is there a reference platform or power budget for "tactical-edge deployment on SWaP constrained platforms"?**
    Answer: The topic does not specify an exact reference hardware model or power budget. Proposers are expected to define their own representative constrained target hardware and demonstrate an information-per-watt improvement relative to baseline systems.

34. **The topic describes Phase II as a 24-month period, while the BAA award table shows an 18-month base and a 6-month option, and several quantitative targets and capstone demos fall in Months 21 and 24. Must the targets be met within the 18-month base, or by the end of the option?**
    Answer: You must plan to achieve and verify the core metric thresholds during the **18-Month Base** period. The Option period is explicitly dedicated to scaling the prototype to the four-modality configuration, integrating with an analytic/data-fabric environment, hardening for tactical-edge SWaP constraints, and hosting the final capstone scenario evaluations.

35. **Will the government provide capstone scenarios, baselines, or the energy-instrumented testbed, or should the proposer supply or propose them?**
    Answer: No, the Government will not supply these upfront. Proposers are expected to supply or propose them.

36. **Will the government provide the SMEs and doctrine-aligned rubrics used to rate the "why / why not" traces, or should we propose them?**
    Answer: Rubrics: Proposed by you, finalized at kickoff. You should outline and propose a doctrine-aligned evaluation rubric in your proposal. The final rubric used by the SMEs to rate trace coherence (targeting ≥ 85%) will be negotiated and finalized with the PM at kickoff. DARPA will coordinate and provide the operational SMEs required for the Month 9 trace-coherence evaluation and the Month 21 commander-intent/audit-interface exercises.

37. **What level of data should we plan for in Phase II: unclassified, CUI, or classified data and environments? Will cleared data or environments be available during execution, given the proposal itself must be unclassified or CUI?**
38. **Answer:**

- **Data Planning Level: Controlled Unclassified Information (CUI) / ITAR.** According to the AMCOMI PSD, this is strictly a **CUI program with explicit ITAR restrictions** The official recommended classification is **CUI.**
- **Cleared Environments:** No classified data or cleared physical environments are projected during execution. All demonstrations, software development, and integrations will occur in restricted, export-controlled (FEDCON unclassified/CUI) environments. Proposers do not need cleared facility access but must possess a **CMMC Level 2 self-assessment** to handle the export-controlled technical data and CUI.

39. For Phase-I-equivalent evidence, how is the at-least-20% missing or delayed-input condition defined: fraction of observations, sensing streams, time steps, or another denominator?

Answer: The topic states that proposers must *"quantify performance under at least 20 percent missing or delayed inputs."* The text does not mandate a single rigid denominator, meaning the 20% can be framed as a fraction of **total observations/events**, **time steps**, or **sensing stream packets** within your shared evaluation task. You should define your chosen denominator clearly in your technical proposal.

40. For the two-modality entry demonstration, how is the difference of at least two orders of magnitude in native tempo measured when a source is event-driven or bursty?

Answer: tempo differences are measured by comparing the **revisit cadences, generation cycles, or mean inter-arrival times** of the sources. For example, the BAA illustrates this by contrasting overhead imagery arriving at revisit cadences of *minutes to hours* with RF/SIGINT resolve times of *seconds to minutes* or maritime AIS updating continuously at *sub-second tempos*.

41. Which unclassified multi-domain data, reference tasks or emitter/power models are available to all offerors for the quantitative entry comparison?

Answer: None. There are no unclassified multi-domain data, reference tasks, or emitter/power models provided by the Government to all offerors for the proposal phase. Offerors must use their own synthetic, simulated, or open-source data (such as public AIS ship-tracking and public satellite imagery) to substantiate the entry criteria.

42. Would a hybrid architecture be in scope if a graph-based or neuro-symbolic layer makes and traces every decision, and a language model handles only the operator interface and translation into that structure?

Answer: Yes, this is in scope. The BAA explicitly rules out *"narrative-only explanations produced by prompted language models"* and *"monolithic end-to-end deep networks without compositional or causal structure"* for core decision-making. However, a hybrid architecture where an inspectable, deterministic graph-based/neuro-symbolic layer makes and traces every decision, and a language model used **only** as a translation interface for human-machine interaction, is consistent with the topic's objectives.

43. For the feasibility comparison against an agentic-LLM baseline, may the offeror's own prior LLM-based system serve as that baseline?
Answer: Yes. The BAA allows the use of any appropriate mainstream baseline family (including agentic-LLM orchestration) on a shared task.