

Decentralized Artificial Intelligence through Controlled Emergence (DICE) Questions and Answers (Q&As)

This document summarizes Q&As related to the DICE program, based on the Broad Agency Announcement, HR0011260010, posted to SAM.gov on June 10, 2026.

1. Are FFRDCs/UARC able to participate as performers?

Please refer to the DICE BAA.

2. What is the nature of agents in the DICE scope? (1) autonomous robotic/embodied AI agents, eg heterogeneous drones that are distributed (i.e. not controlled centrally) and can form ad hoc teams to execute missions in dynamic physical environments (albeit in simulation for the purposes of DICE) that need to be safe, robust or (2) robotic agents as in (1) that are orchestrated centrally, or (3) "agentic" AI systems, such as foundational models that operate on the web, and orchestrated centrally to make composite systems to execute user tasks, eg make my travel arrangements of hotel, airfare, car hiring, entertainment etc.?

Please refer to the DICE BAA.

3. What is the envisioned user interaction and control?

The envisioned user interaction is primarily focused on high-level mission assignments, with the system of agents then autonomously managing the execution details. Please refer to the DICE BAA for additional details.

4. How does DARPA intend to evaluate performer claims around “controlled emergence,” particularly with respect to measurable improvements in scalability, adaptability, and resilience under degraded, adversarial, or uncertain operating conditions? Are there preferred approaches for demonstrating that emergent behaviors remain bounded, observable, and operationally trustworthy?

DARPA intends to evaluate "controlled emergence" through a structured, metric-driven process managed by the Technical Area 3 (TA3) Test and Evaluation (T&E) team. The evaluation is primarily based on the measurable outcomes of the collective system's performance. Please refer to the DICE BAA for additional details.

5. Does DARPA anticipate providing common simulation environments for DICE evaluation, or will performers primarily develop their own? If performer-developed environments are expected, what level of fidelity and validation credibility would be considered sufficient for demonstrating safe autonomy and robustness across multi-step inference and adversarial conditions?

DARPA will provide a common simulation environment. This environment will be developed and managed by the dedicated TA3 performer(s), not the TA1/TA2 research performers. Please refer to the DICE BAA for additional details.

6. Given DICE’s emphasis on robust-to-compromise architectures and adversarial resilience, does DARPA anticipate alignment with existing federal AI assurance or security

DICE Q&A

frameworks (for example, NIST AI RMF or Zero Trust concepts), or should performers treat assurance methodologies as an open research dimension within the program?

Resilience is a key component of the research challenge being addressed by DICE and we anticipate continued alignment with existing frameworks.

7. The DICE framing contrasts decentralized self-organizing collectives with centralized orchestration approaches. Does DARPA anticipate encouraging multi-organization performer teams that combine complementary capabilities across autonomy, formal assurance, and adaptive AI methods? Additionally, are hybrid approaches combining statistical ML with symbolic or formally verifiable reasoning viewed as within scope for the program vision?

Both multi-organization teams and hybrid technical approaches are acceptable as long as they meet program objectives.

8. Would a robotics perspective be within scope?

Yes, a robotics perspective is within scope. Experiments with physical platforms in real world are not in scope.

9. Can PIs appear on multiple proposals? If you are PI on one submission, can you serve as co-PI on another?

Yes, a PI on one proposal can appear as PI/co-PI on another. However, a PI/Co-PI cannot perform with more than one awardee. The organizations/PIs selected to perform in TA3 cannot be part of any TA1-TA2 team in the program.

Note that a planned downselect is scheduled for the end of Phase 2 in TA1-TA2.

10. When the solicitation talks about "simple local rules" yielding "robust global behavior," is DICE scoped to ML-policy-based agent rules specifically, or does it admit physics-based deterministic propagators (e.g., coupled-oscillator dynamics, field-based coupling) as a valid local-rule substrate?

DICE is not scoped to only ML-policy-based agent rules. The BAA explicitly encourages diverse approaches drawn from a wide range of scientific disciplines. Please refer to the DICE BAA for additional details.

11. Does "robust global behavior" require strict determinism (same initial conditions -> same trajectory), or does DICE expect stochastic convergence with statistical guarantees? What bound does DICE have in mind for "controlled" in "controlled emergence"?

DICE expects stochastic convergence with statistical guarantees, not strict determinism. The notion of "controlled" is about maintaining mission alignment and role coherence within certain bounds, not always forcing a single, predictable outcome. Please refer to the DICE BAA for additional details.

12. Are proposers expected to develop the coordination substrate from scratch within DICE, or can proposers bring an existing substrate (internal-R&D technology base, prior work) and use DICE funding to advance specific DICE-program properties on top?

DICE Q&A

Please refer to the DICE BAA.

13. When DICE talks about "robust to failure or compromise of individual agents" and resilience against "rogue" AI agents -- is the resilience expectation policy-level (the collective routes around a compromised agent), substrate-level (the coupling rules tolerate adversarial inputs), or instrumentation-level (the measurement layer detects and adapts)?

Resilience at all levels is in scope. Please refer to the DICE BAA.

14. Are we considering a fully decentralized setting or can we consider the possibility of having a high-level coordinator, e.g., a (autonomous) commander who determines subtask delegations of some form, either statically at the beginning of the task execution or dynamically at runtime?

Please refer to the DICE BAA.

15. Does the program consider any specific adversarial models?

Yes, the program considers several specific, albeit broad, adversarial models focused on manipulating information and behavior. Please refer to the DICE BAA for additional details.

16. Will government performers be selected for TEV&V (testing, evaluation, validation and verification) activities of the technology developed by the performers. If yes, how could government performers get involved?

Government performers are expected to assist with TEV&V through complementary activities, such as, helping identify DoW-relevance of use cases and suggest any needed refinements. Interested potential Government performers are encouraged to contact DARPA directly.

17. Is evaluation of proposed agent coordination/local inference techniques on physical autonomy platforms (e.g., UAVs/UGVs/USVs) within the scope of the proposed activities and will justified budget for hardware (e.g., robots) acquisition be allowed?

No, evaluation on physical platforms and budgets for hardware acquisition are out of scope for this program.

18. Are there any restrictions on source/country of origin of open-source and third-party software that might be used in the project?

There are no additional program-specific restrictions on software used in DICE beyond the generally applicable rules and regulations.

19. Does the research involve Controlled Unclassified Information (CUI), or is it fundamental research?

Please refer to the DICE BAA.

20. Does the project require US citizenship/Permanent residents (for senior personnel and/or students)?

Please refer to the DICE BAA.

DICE Q&A

- 21. What is the funding size?**
Funding details will not be made available.
- 22. How many phases does the project have? How will the project be evaluated?**
Please refer to the DICE BAA.
- 23. Do all performers need to be US citizens/persons or only PIs/Co-PIs? For example, can I hire student research assistants who are foreign nationals studying at US universities on an F1 visa? Can non-US citizens working in US companies be able to work on DICE?**
Please refer to the DICE BAA.
- 24. For DICE, is it in scope to adapt (fine-tune) AI models to improve performance in distributed agent coordination scenarios?**
Diverse approaches that meet program objectives are encouraged. TA1 coordination methods need to support heterogeneous agents provided by TA3. Please refer to the DICE BAA.
- 25. Is it in scope to adapt / fine-tune AI models to better control scheming behaviors, strategic deception and misalignment? And improve relevant observability?**
Diverse approaches that meet program objectives are encouraged. TA2 control methods need to support heterogeneous agents provided by TA3. Please refer to the DICE BAA.
- 26. For Phase I, should performers expect evaluations to focus primarily on benign decentralized coordination, or should we anticipate adversarial conditions such as compromised agents, coordination poisoning, indirect prompt injection, Byzantine-style degradation, or role-coherence failure across the collective?**
Phase 1 evaluations will focus primarily on benign decentralization, with adversarial robustness being the main focus of Phase 2. Please refer to the DICE BAA for additional details.
- 27. Will DICE distinguish between decentralized orchestration and decentralized governance? Specifically, will performer systems be evaluated only on task coordination, or also on their ability to preserve authority boundaries, trust lineage, and recoverability when local agents become unreliable or misaligned?**
Please refer to the DICE BAA.
- 28. For decentralized collectives that use adaptive or self-improving coordination mechanisms, should performers assume that governance mechanisms themselves may be treated as targets of optimization, drift, or adversarial manipulation during evaluation?**
Proposals can describe assumptions needed (including supported attack models) by their technical approach.
- 29. Is it still within the program's scope if an agent spontaneously takes on the role of a coordinator resulting in a system that resembles a centrally coordinated multi-agent system?**

DICE Q&A

Diverse approaches that meet program objectives are encouraged. Please refer to the DICE BAA for additional details.

30. Can a team propose to more than one TAs, e.g., TA1+TA2, given that a method that addresses the program's challenge may require intrinsic coordination of self-organization and local inference control?

The BAA explicitly states that a single proposal must cover both TA1 and TA2.

31. DICE appears to avoid centralized orchestration while still recognizing the need to fuse partial, changing, and potentially conflicting information across an agent collective. Would DARPA consider a non-orchestrating capability-state fusion layer to be in scope, where agents or collectives can discover the next relevant or eligible capability as mission conditions evolve, without centrally assigning tasks or hardcoding workflows? If so, would DARPA encourage proposers with this kind of complementary capability to team with organizations focused on decentralized coordination and local inference control?

Yes and Yes.

32. Will GFE be provided for representative autonomous platforms, algorithms, or interfaces?

No. DICE will use simulation environments, and no physical platforms will be provided.

33. During the contracts intro presentation, the presenter commented that the program could be non-fundamental. However, during the program manager presentation, “fundamental research” was mentioned several times. Can you clarify the fundamental vs non-fundamental research desired on the program again as it relates to the overall program and individual TAs?

Please refer to the DICE BAA.

34. Should the application in the proposals be directly related to defense environment and war agents or it could be indirectly related such as coordination across entities of a micro-grid?

The use case must be DoW-relevant (Department of War-relevant). Please refer to the DICE BAA for additional details.

35. Will there be a formal RFA document explaining the requirements, timeline, etc.?

The BAA is the formal solicitation document.

36. Is there any interest in Human in the loop decentralized systems? E.g. it seems that particularly for DoW there is a need to consider that some agents might be human actors (e.g., soldiers and other decision markers) coordinating with AI agents.

Diverse approaches that meet program objectives are encouraged. TA3 simulation environment could include simulating human agents. Please refer to the DICE BAA.

37. Coordination relying solely on peer-to-peer coordination can place a huge burden on communication in large-scale multi-agent systems. Is hybrid, locally centralized but globally decentralized systems in scope?

DICE Q&A

Diverse approaches that meet program objectives are encouraged. Please refer to the DICE BAA.

38. Classical mechanism design often assumes a mechanism can aggregate information from all agents through a central auctioneer or planner. Is this at odds with the requirement of a fully decentralized setting or using purely peer-to-peer communication and coordination?

Diverse approaches that meet program objectives are encouraged. Please refer to the DICE BAA.

39. For peer-to-peer coordination how is DICE establishing "scope" of agent communication? ie -> text / visual / touch / other sensory methods?

Diverse approaches that meet program objectives are encouraged. Proposals can describe assumptions on their technical approach, including supported modalities such as text and vision. Please refer to the DICE BAA.

40. Are multiple awards to TA3 anticipated, or a single award?

Please refer to the DICE BAA.

41. Should time constraints on decisions be considered as a criteria for the program. It is important to provide these constraints to ensure that the systems can apply to real world planning problems.

Please refer to the DICE BAA.

42. Will the program provide LLM access to frontier models as that will be important for performers and can be a big source of cost.

Proposals should describe the team's existing compute capability, access to compute such as (but not limited to) national initiatives, and budget additional compute cost without assuming any GFE. For closed models, the program expected to heavily use versions that are not the most recent release because a key goal of the program is to support model heterogeneity. For open-weight models, the proposers are strongly encouraged to review the state of art in efficient agentic inference frameworks and the rapid recent progress in small reasoning models when estimating compute cost. Please refer to the DICE BAA and the templates for abstract and proposal on budgeting cost for compute.

The BAA supports multiple award instruments to support diverse performers with different teaming arrangements.

43. Will there be constraints placed on the LLM number of parameters?

There are no constraints on the number of parameters for the models used.

44. Could you please provide a couple examples of relevant scenarios that would be simulated for the experiments?

Please refer to the DICE BAA.

45. Could you please explain the types of heterogeneity you anticipate within the AI collectives under consideration?

DICE Q&A

Please refer to the DICE BAA.

46. Can a company bid on both TA-1/ TA-2 and TA-3 and then only work on the one selected?

Yes.

47. By AI agents - are we limited to only LLM? That is no AI such as RL and programs of robotics?

No, the program is not limited to only LLMs, for example, agents proposed in TA3 could use VLMs or VLAs. Program aims to build an adaptor that supports heterogeneous agents. The proposals can describe assumptions on the adaptor interface. Please refer to the DICE BAA for additional details.

48. Do the proposals need to be written in 12 pt? (both abstract and proposals)?

Please refer to the DICE BAA.

49. Are remote attestation protocols in scope?

Diverse approaches that meet program objectives are encouraged. Proposals to a technical area are encouraged to address all the key technical challenges in that TA. Please refer to the DICE BAA.

50. Over what time horizon must alignment be maintained? Hours, days, weeks? How many inference steps does this translate to in your operational concept?

The BAA specifies the time horizon in terms of inference steps, not clock time. The goal is to sustain alignment over increasingly long and complex missions. Please refer to the DICE BAA for additional details.

51. What is the definition of an agent? To what extent agents have knowledge of each other? What is the bandwidth of communication? How do they communicate? Is correctness of any interest? Will DARPA give us data for evaluations?

Please refer to the DICE BAA.

52. Are there specific use cases in mind? Like War-gaming, logistics, etc?

Please refer to the DICE BAA.

53. In the context of adversarial attacks, do we expect robustness in terms of attacks on the network or on the local model (or both)?

The focus is on attacks targeting the local models and the information they receive, not attacks on the underlying network infrastructure. Please refer to the DICE BAA for additional details.

54. Is theoretical research on Emergent behavior in physical systems in scope?

Diverse approaches that meet program objectives are encouraged. Proposals to a technical area are encouraged to address all the key technical challenges in that TA. Please refer to the DICE BAA.

DICE Q&A

55. Will TA3 approaches require generalization, or can they be use case specific?

The TA3 approach itself must be general enough to support the program's goals, but it will be implemented through specific, DoW-relevant use cases. Please refer to the DICE BAA for additional details.

56. Will compute for TA3 be available as GFE?

Proposals should describe the team's existing compute capability, access to compute such as (but not limited to) national initiatives, and budget additional compute cost without assuming any GFE. For closed models, the program expected to heavily use versions that are not the most recent release because a key goal of the program is to support model heterogeneity. For open-weight models, the proposers are strongly encouraged to review the state of art in efficient agentic inference frameworks and the rapid recent progress in small reasoning models when estimating compute cost. Please refer to the DICE BAA and the templates for abstract and proposal on budgeting cost for compute.

The BAA supports multiple award instruments to support diverse performers with different teaming arrangements.

57. Are subcontractors required to have a CAGE code?

Please refer to the DICE BAA.

58. When would you aim to provide abstract feedback?

Within 4 weeks of abstract submission.

59. Can an org sub on a TA1/TA2 proposal if they're priming TA3 (with appropriate firewalls in place)?

No.

60. Team v team in phase 2 evaluation—could you define teams better? Two teams where both can achieve objective or one against the other? Meaning, zero-sum or not?

Please refer to the DICE BAA for additional details about the ideal simulation environment and use cases.

61. If a prime bid a TA1/TA2 combined proposal, would you award a partial TA1 or TA2?

It is highly unlikely. The BAA is structured to treat TA1 and TA2 as a single, integrated effort.

62. Can a single proposer team include multiple model/agent vendors as subs or teammates, given the program's emphasis on heterogeneous and plug-and-play agents?

Yes.

63. Can a single TA1 company be on multiple proposals?

Yes, a single TA1 company can be on multiple proposals. The organizations selected to perform in TA3 cannot be part of any TA1-TA2 team in the program. An organization can be a part of multiple TA1-TA2 teams.

DICE Q&A

Note that a planned downselect is scheduled for the end of Phase 2 in TA1-TA2.

64. If TA1 + TA2 must be proposed together, can one team contribute mainly the infrastructure / coordination layer while different partners provide the agent/control pieces?

Yes.

65. Does DARPA want the local adapter to be vendor-agnostic by design, and would they view multi-vendor interoperability as a strength in evaluation?

Yes.

66. For government purpose-rights requirements, how should teams handle cases where closed-model vendors restrict access to internals, weights, or source code?

Please refer to the BAA for example technical approaches for closed-model agents.

67. In evaluation, will performers be expected to demonstrate portability across multiple underlying agent architectures/vendors, or is one vendor enough in Phase 1?

Proposals should describe assumptions on their technical approaches. The program aims at supporting heterogeneous agents.

68. Can the same company participate on multiple teams aligned to different technical approaches or different underlying vendors?

The same company can be on multiple proposals. The organizations selected to perform in TA3 cannot be part of any TA1-TA2 team in the program. The same organization can be a part of multiple TA1-TA2 teams.

Note that there is a planned downselect at the end of Phase 2 in TA1-TA2.

69. What level of access to the agent is assumed for TA2 control in closed-model settings: API only, logits, traces, tool calls, or something deeper?

For closed-model settings, the BAA assumes access is limited to what is available through standard agentic frameworks, not deep internal states. Proposals should describe assumptions needed for their technical approach. Please refer to the DICE BAA for additional details.

70. For TA1, how much of the score will depend on distributed coordination under degraded comms, versus purely simulated agent reasoning quality?

Please refer to the DICE BAA.

71. I understand TA3 performers cannot work on TA1/TA2. Can a team apply to both though and be selected as a performer on one or the other?

Yes, teams can apply to both, but TA3 performers cannot work on TA1/TA2. Please refer to the DICE BAA.

72. Are all AI agents supposed to be able to communicate with each other?

DICE Q&A

No. The BAA explicitly assumes a sparse communication model where agents do not communicate with everyone.

73. Are all agents aware of the capabilities of other agents?

No. Agents have partial observability and are not aware of the full state of the collective.

74. What is the expected TRL at program completion? Is this a science program (TRL 2–3) or a technology demonstration program (TRL 4–5)?

This program starts as a fundamental science program.

75. Is LLM inference in scope for each agent in this environment?

Yes. The use of modern AI models such as (but not limited to) LLMs, VLMs, and VLAs as the "brains" of each agent is a foundational premise.

76. Is the assumption that individual agents in the simulated or real-world environment will have access to a frontier inference model, but are still unable to communicate directly with peer agents on the far side of the network?

The agents in DICE are expected to be heterogeneous. The architecture of the individual agents and the inter-agent communication constraints are not directly coupled and will be determined by the use case. Please refer to the DICE BAA.

77. Regarding the numbers of agents mentioned across the three phases: do these numbers correspond to the evaluation settings? In other words, will proposed methods be evaluated under scenarios involving those specific numbers of agents?

Yes, those numbers directly correspond to the scale of the formal evaluation settings in each phase.

78. Regarding team size: do you have any guidance on the preferred size of proposal teams? More specifically, how small can a team be, and are smaller PI groups viewed as less competitive?

The BAA provides no guidance on teaming. Diverse approaches that meet program objectives are encouraged. Proposals should describe assumptions needed for their technical approach.

79. Since DICE focuses on local control and peer-to-peer coordination, we wonder if the AI agents themselves will be provided or need to be built by the proposing team. What type of agents do you expect to use, LLMs, VLMs, RL agents, or DNN agents?

The AI agents themselves will be provided by the TA3 T&E performer, not built by the TA1/TA2 research teams. The BAA mentions a variety of modern models, including LLMs, VLMs, and VLAs. Proposals should describe assumptions needed for their technical approach.

80. Can you provide further clarifications on the attack model with respect to the resilience objectives? For example, shall we consider advanced attacks that learn, model, and stimulate agents?

Diverse approaches that meet program objectives are encouraged. Proposals to a technical area are encouraged to address all the key technical challenges in that TA.

DICE Q&A

Proposals should describe assumptions needed for their technical approach, including description of the attack model. Please refer to the DICE BAA.

81. Is the program about embodied AI agents? I.e., is the goal here to coordinate a team of physical agents or a team of virtual agents?

Please refer to the DICE BAA.

82. I understand that in the long term the goal is to coordinate heterogeneous teams. However, even the coordination of homogeneous agents is a challenge today. Is it therefore reasonable to assume that agents are homogeneous? At least for Phase I and II.

Diverse approaches that meet program objectives are encouraged. Proposals to a technical area are encouraged to address all the key technical challenges in that TA. Proposals should describe assumptions needed for their technical approach, including any assumption on the nature of agents.

83. Can you give an example of a task that we want the team to solve? It may provide useful guidance on proposed approaches.

Please refer to the DICE BAA.

84. What is the threat model in terms of adversarial capabilities and interactions? Can we assume a specific one?

Please refer to the DICE BAA.

85. What are the input modalities of each agent? Can we choose it? Does it have to be generic?

The input modalities will be defined by the TA3 performer as part of the simulation environment. Proposals should describe assumptions needed for their technical approach, including supported modalities.

86. What kind of agent should we consider? Is it only LLM? Can it be generic? What about tasks?

Please refer to the DICE BAA. Proposals should describe assumptions needed for their technical approach.

87. Is the size of the interaction between agent important? What about robustness of the encoded data, and the time the communication channel is being used?

Diverse approaches that meet program objectives are encouraged. Proposals to a technical area are encouraged to address all the key technical challenges in that TA. Please refer to the DICE BAA.

88. What will be the proposal format requested by DICE? Will it be a traditional written proposal (vs video or in-person presentations)?

Please refer to the DICE BAA.

89. Given the scale in terms of number of agents and interactions envisioned for this program, will hardware infrastructure be provided by DARPA to support this?

DICE Q&A

Proposals should describe the team's existing compute capability, access to compute such as (but not limited to) national initiatives, and budget additional compute cost without assuming any GFE. For closed models, the program expected to heavily use versions that are not the most recent release because a key goal of the program is to support model heterogeneity. For open-weight models, the proposers are strongly encouraged to review the state of art in efficient agentic inference frameworks and the rapid recent progress in small reasoning models when estimating compute cost. Please refer to the DICE BAA and the templates for abstract and proposal on budgeting cost for compute. The BAA supports multiple award instruments to support diverse performers with different teaming arrangements.

90. Should human input be considered during the decentralized planning phase beyond specifying the initial goal?

The goal of the program is to develop an autonomous collective that operates without continuous human intervention after the initial mission is assigned by the human user. Please refer to the DICE BAA for additional details.

91. Can you clarify the timeline for TA3? When should the simulation environment be ready for TA1&2 evaluation? Also, can TA3 simulation environment be based on current open-source simulators?

Yes, the BAA provides a clear timeline and encourages the use of open-source tools.

92. Would you consider the approaches (models, ideas, solutions) to make the agents have “complementary” representations (as opposed to “similar” representations) within scope? Example is that the agent A builds a representation X and agent B builds representation Y, and these are not similar ($X \neq Y$), but jointly coherent ($X+Y = \text{world}$)

Yes.

93. The centralized teams perform better when the missions (tasks, workflows) are well understood. Given that, what is your view on the hybrid team organizations – i.e., a combination between local (decentralized) and global (partially centralized) controls where agents assume not only the “task roles” but also “organizational roles” (e.g., agent A assigns tasks to “subordinate” agents B, C, and D and/or ensures the coordination between them), while changing / adapting these over time as mission evolves?

Please refer to the DICE BAA.

Would you consider the self-assessment of the agents about the representations of other agents (and/or about the uncertainties of other agents) that they depend on – to be within scope?

Yes.

94. Would agents be learning about other agents, and if so, would concepts of imitation or “learning / assembling the behaviors from other agents” be within scope?

Yes. Diverse approaches that meet program objectives are encouraged. Proposals to a technical area are encouraged to address all the key technical challenges in that TA.

DICE Q&A

Proposals should describe assumptions needed for their technical approach. Please refer to the DICE BAA for more details.

- 95. Should we assume that agents' skills are "fixed", i.e., that underlying "agent models" (LLMs, foundational cross-modal models, NNs) are fixed? Would there still be an option to "contextualize" the agent's behavior by the knowledge of the agent about other agents (and not just the tasks)**

The underlying models (LLMs, foundational cross-modal models, NNs) are fixed. Agents can learn, for example, via memory that can be controlled by the adaptor. Diverse approaches that meet program objectives are encouraged. Proposals should describe assumptions needed for their technical approach.

- 96. Will DICE utilize existing DoW HPC resources for offerers to use or will proposals need to address HPC needs? If using existing, please provide details for our planning purposes.**

Proposals should describe the team's existing compute capability, access to compute such as (but not limited to) national initiatives, and budget additional compute cost without assuming any GFE. For closed models, the program expected to heavily use versions that are not the most recent release because a key goal of the program is to support model heterogeneity. For open-weight models, the proposers are strongly encouraged to review the state of art in efficient agentic inference frameworks and the rapid recent progress in small reasoning models when estimating compute cost. Please refer to the DICE BAA and the templates for abstract and proposal on budgeting cost for compute. The BAA supports multiple award instruments to support diverse performers with different teaming arrangements.

- 97. Does DICE want to consider unreliable communications, comms-limited or comms-denied environments?**

Yes.

- 98. Some tasks are best solved by multiple agents self-selecting out of a swarm and coordinating together. Does DICE want to see coordinated sub tasking?**

This is in scope.

DICE Q&A

- 99. When you say a “low-fidelity simulation environment,” do you mean environments such as Isaac Sim, or reduced-order models created using synthetic data generated from finite element models?**

Diverse approaches that meet program objectives are encouraged. Proposals should describe assumptions needed for their technical approach. Please refer to the DICE BAA.

- 100. For TAs, Phase 1 focuses on Physical+Social and Phases 2 and 3 focus on Physical+Social+Cyber. What does each of Physical, Social, and Cyber mean in this context? Can you be more specific?**

Please refer to the DICE BAA.

- 101. Related to the simulation domains, is “social” (on TA3 slide but not indicated on the schedule chart) considered an independent domain or considered part of both physical and cyber simulations?**

Please refer to the DICE BAA.

- 102. Can you please clarify how you envision the evolution of different simulation domains from phase 2 to phase 3? For example, on the schedule chart, is the “Cyberphysical simulation” in phase 3 keeping the two domains (e.g., Minecraft for physical and CAGE4 for cyber) separate, or combined somehow?**

Please refer to the DICE BAA.

- 103. Are there example and papers relevant to TA3 which will help us understand better what DICE is looking for?**

Yes, the BAA provides a list of citations that are highly relevant to the concepts for TA3.

- 104. Are industry-academia collaborations required?**

No.

- 105. Will a list of attendees be made public to encourage collaboration?**

Yes.

- 106. Would Darpa allocate funding and/or resources for computation? What would be the amount, since the large number of LLM agents would require large amounts of compute.**

Proposals should describe the team’s existing compute capability, access to compute such as (but not limited to) national initiatives, and budget additional compute cost without assuming any GFE. For closed models, the program expected to heavily use versions that are not the most recent release because a key goal of the program is to support model heterogeneity. For open-weight models, the proposers are strongly encouraged to review the state of art in efficient agentic inference frameworks and the rapid recent progress in small reasoning models when estimating compute cost. Please refer to the DICE BAA and the templates for abstract and proposal on budgeting cost for compute. The BAA supports multiple award instruments to support diverse performers with different teaming arrangements.

107. What modalities do you envision/require for the foundational models/agents? LLMs, VLMs, MMLMs?

The BAA envisions a mix of modalities. Please refer to the DICE BAA for additional details.

108. I assume that the agents collaborate, communicate and interact to plan and also execute actions in the simulated world for the given mission. Is this correct, i.e both planning and execution are required?

Yes.

109. What scenarios are in scope, e.g. scenarios with a single overall goal, or requirement for multiple (sub) goals?

Please refer to the DICE BAA.

110. Is the whole multi-agent system envisioned to collaborate as a single group in executing the mission or is it required to have collaborating subgroups?

Please refer to the DICE BAA.

111. Are DICE distributed agents foundational pretrained models, eg LLMs, VLMs or agents that are distributed but perform based on analytical methods, e.g. auctions, swarm consensus algorithms, control theoretic techniques, Barrier Functions?

Please refer to the DICE BAA.

112. Are agents engaged in planning only or planning and execution in simulations?

Please refer to the DICE BAA.

113. Could DICE agents be mixtures of foundational pretrained models, and agents whose individual reasoning is based on analytic approaches and coordinate based on techniques such as auctions?

Yes, diverse approaches that meet program objectives are encouraged. Please refer to the DICE BAA for additional details.

114. What sensors are these agents supposed to have e.g. vision, actuation?

The agents' sensors and actuators will be abstract and defined by the simulation environment. Please refer to the DICE BAA for additional details.

115. Is there a restriction on how many proposals one can serve as PI and/or co-PI?

There is no restriction.

116. What is meant by "interactions"?

The BAA explicitly defines "interactions" as "messages exchanged between the agents". This is one of the notional metrics for mission complexity described in the BAA. TA3 proposals are expected to expand on the notional metrics with use-case specific metrics.

117. Does DARPA have an estimated maximum wall-clock time for the competition/evaluations during a phase? For example, if the competition/evaluation took about a day per TA1/TA2 team, would that be too long?

The evaluation timeline will be determined by the TA3 performer, who is responsible for designing the T&E platform. DARPA expects the TA3 performer to engineer a solution that makes running these evaluations practical within a reasonable timeframe (likely hours or days, not weeks).

118. Is it the case that TA3 will provide a way to assess the TA1/TA2 teams' performance across a range of conditions (such as increasing task difficulty, increasing communications costs, increasing rate of environmental change or surprise)? Is it the case that the team vs team competition is also envisioned to be carried out across a range of conditions, or just at a single representative set of conditions?

Yes, the TA3 performer is responsible for creating a rich evaluation environment that tests performance across a range of conditions. The BAA tasks TA3 with developing use cases that: "support missions of different complexity"; are set in a "non-stationary environment to test adaptability"; "elicit emergent collaboration, competition and collusion behaviors," which implies testing against cascading perturbations. The evaluations, including the team vs team competition, are not meant to be a single static test but a comprehensive assessment of the system's robustness and adaptability under varied and challenging conditions.

119. Will TA3 develop the API between the local adapter and the simulation environment, as well as the API between the local adapter and all agents?

Yes. The TA3 proposal is expected to address all aspects of creating the complete T&E platform. The TA1/TA2 "local adaptor" is a component that plugs into this platform. The BAA states that TA3 provides the "simulation environment... including the AI agents to which TA1-TA2 adaptors will connect". The BAA also notes that early in the program, "DICE will converge to a standard adaptor interface", which will be managed by TA3 and DARPA.

120. By social reasoning/social domain, does that mean the simulated scenario will involve simulated human behavior in the world as part of the environment the agents must operate in, with agents trying to anticipate and shape the reactions of simulated humans? Or does it mean there is a social element in the agent-to-agent interactions needed to complete the task/mission, such as agents maintaining trust or having preferred relationships with other specific agents?

It means the latter. The "social" domain refers to the agent-to-agent interactions. The program focuses on AI-to-AI coordination. The BAA's description of social reasoning is tied to concepts like: agents engaging in "negotiations" or seeking "help from other agents"; developing a "distributed reputation management approach"; and handling

DICE Q&A

"colluding compromised agents" and deceptive messages. These are all social dynamics within the agent collective, not interactions with simulated human populations

121. By cyber reasoning/cyber domain, does that mean the simulated scenario will involve simulated cyber attacks/cyber defense? Or that it involves a simulated cyberphysical system, such as a scenario where simulated sensors and electronic communications are used to control mobile systems in simulated physical space? Or something else?

The BAA clarifies this distinction:

In Scope: A simulated cyberphysical system. The overall scenario involves agents controlling assets (like drones) in a physical space, using information from simulated sensors and communicating electronically. The "cyber" domain here refers to the information and communication layer that these agents operate on.

Out of Scope: Pure cybersecurity scenarios. The BAA explicitly states, "Use-cases that exclusively focus on cybersecurity such as network emulation/simulation frameworks are not in scope".

So, "cyber" reasoning in DICE is about managing and protecting the information and communication systems that enable the agents to coordinate and control their assets in the simulated physical world.

122. Could you please confirm that this a CPFF type proposal is acceptable and, if there are any additional requirements we should address, advise us accordingly?

Section 1, "Overview Information," lists the "Types of Instruments that may be Awarded" as Procurement Contract, Other Transaction (OT) for Prototype Agreement, OT for Research Agreement, and Cooperative Agreement. A CPFF contract falls under the general category of a Procurement Contract.

123. I am contacting you in regard to the DICE solicitation #HR001126S0010. My question concerns the resource sharing requirement if we are pursuing an OT for Research Agreement. I found the Guide to Research Other Transactions Under 10 U.S.C 4021 very helpful in explaining the requirement, however, we were hoping you could provide more insight into the amount of resource-sharing that is expected for this particular solicitation.

As stated in 10 USC § 4021(e)(2),

(2) to the extent that the Secretary determines practicable, the funds provided by the Government under a cooperative agreement containing a clause under subsection (d) or a transaction authorized by subsection (a) do not exceed the total amount provided by other parties to the cooperative agreement or other transaction.

No other determination has been made at this time.

124. I have downloaded the documents and am filling them, but I noticed the Apply button on the Grants.gov webpage is inactive (see attached screenshot). I was wondering if there is a date when the applications will open. Could you please let me know when that date is?

All abstracts and proposals must be submitted to the DARPA Broad Agency Announcement Tool (BAAT). See BAA sections 5.1 and 6. Please visit [Proposer Instructions: General Terms and Conditions](#) for specific information regarding submission methods through the Broad Agency Announcement Tool (BAAT).

125. Is it allowable for an underlying algorithm to know the adjacency and communication connectivity graph of all agents?

It is unlikely that an algorithm would be able to maintain the complete, up-to-date communication graph of all agents. This would violate the program's core principles of decentralization and partial observability. Therefore, any solution requiring a centralized repository of the full, real-time communication graph is considered out of scope.

126. Please advise if travel for domestic and/or international conferences is allowed.

The program encourages the open publication of research results. Section 2.8 ("Intellectual Property and Data Rights") states, "DARPA expects most research accomplished under DICE to be open-sourced in the form of publicly posted... or published papers and source code." Attending domestic and international conferences is a primary and standard method for disseminating such research to the scientific community. Therefore, budgeting for travel to present DICE-related papers at relevant, reputable conferences may be considered justifiable expense.

127. A follow-up question on the program metrics, building on published Q&A #116, which defines "interactions" as "messages exchanged between the agents." For an approach in which agents coordinate primarily by reading and writing a shared environment (stigmergic coordination) rather than by exchanging direct agent-to-agent messages, what is the countable unit of "interaction" under the Section 2.5 / Table 2 metrics (mission complexity, scalability, time-to-recover)? Specifically:

1. Is each agent read from, or write to, the shared environment counted as one interaction — or are environment-mediated approaches scored as effectively O(1) message complexity because no direct messages are exchanged?

2. If environment reads/writes are counted, are sensing-coupled writes (where the signal an agent deposits is physically tied to work it performs, and costs the agent proportionally) distinguished from assertion writes (free, unverified broadcasts)? The distinction is

material to the resilience metric: the two classes have opposite economics under Byzantine or deceptive failure, and an evaluation that counts them identically cannot see that difference.

3. Where the BAA notes the notional metrics will be refined by TA3 into use-case-specific metrics, will message-based and environment-mediated approaches be evaluated on a common scale — e.g., total coordination-relevant information transferred, or coordination cost per unit of mission performance — so that neither class is structurally advantaged by the choice of countable unit?

We ask because a literal reading of Q116 would score a low-message, environment-mediated system as either trivially optimal (zero messages) or entirely unmeasured — both misleading. A clarification of the countable unit would let proposers in this class present scalability and resilience numbers commensurable with message-passing approaches.

The BAA notes that the metrics provided are "generic and notional" and that the TA3 performer is explicitly tasked with refining and extending them. Your question highlights exactly the kind of refinement that would be necessary to fairly evaluate a stigmergic coordination approach against a message-passing one. A proposal using a stigmergic approach should propose its own justifiable definition of an "interaction" as a starting point for this refinement process. Such coordination and control approaches will tradeoff other metrics of interest in DICE such as scalability and resilience. Proposed approaches need to consider all DICE-relevant metrics.

128. The proposal template lists Section 7 Statement of Work (which is not part of the page count) before Section 8 Schedule and Milestones (which is part of the page count). Would it be acceptable to instead put the Statement of Work directly after the Schedule and Milestones, so that all the parts of Volume 1 that are included in the page count limit are contiguous?

It is strongly recommended that you do not reorder the sections and that you adhere strictly to the format provided in the solicitation.

129. We are preparing a full TA3 proposal in response to HR001126S0010. We plan to submit an unclassified proposal containing no CUI. We would appreciate clarification on two points:

1. For a non-CUI TA3 proposal, is a concrete, budgeted plan to establish the required CUI environment and CMMC status before award sufficient for the mandatory conformance review, or must either already be in place at the time of proposal submission?

2. Under the current suspension of CMMC Phase II, does Conditional Level 2 (Self), with

DICE Q&A

a current SPRS affirmation, satisfy TA3 award eligibility, or is Final Level 2 (Self) required?

1. Proposers must demonstrate compliance with the applicable CMMC level at the time of contract award and maintain that level throughout contract performance.
2. To be eligible for an award under this solicitation, the CAGE code associated with the proposal must have a valid, unexpired Final Level 1 for TA1/TA2 and Level 2 for TA3 CMMC status recorded in the Supplier Performance Risk System (SPRS). Additionally, the proposer must have a current annual CMMC affirmation submitted by an authorized Affirming Official and recorded in SPRS.

130. On pages 29-30 of the BAA, it asks organizations to include red-lined edits on a Model Agreement with their proposal submission. However, the link for Cooperative Agreements does not lead to a sample agreement. Can you please provide one or a link to where this can be found?

Sample OT Cooperate Agreements can be found

131. Regarding the P4_-_DARPA_Standard_Cost_Proposal_Spreadsheet, the template only references the TAs in the Milestone tab. The Base, OP1, etc. appear to be for a single TA. Can we use the Multi-TA / Multi-Phase version that we have from DARPA (dated Nov 2021), or must we use the supplied P4 and submit a separate document for each Phase / TA combination?

Additional templates for the pricing worksheet can be found at the Proposer Instructions website, which was incorporated into the BAA at Section 4. SUBMISSION INFORMATION, as Proposer Instructions: General Terms and Conditions: <https://www.darpa.mil/work-with-us/proposer-instructions>. See the section titled Cost Proposal Spreadsheet instructions.

132. Could you clarify in this context what are the definitions of a "use case" and a "scenario"? Is a "use case" a broad type of mission or operation, such as "search and rescue behind enemy lines" or "logistics planning for materiel and fuel transport"? Is a scenario the set of starting conditions, environmental injects, and other simulation configuration settings needed to run a specific instantiation of a specific situation that could happen in the use case? Is it expected that a use case can have multiple scenarios, or is there any other hierarchical relationship between the two things?

DICE Q&A

Your description is aligned with the intent of the BAA. A use case is the broad operational context, and a scenario is a specific, configured instance of a situation within that context, designed to produce a measurable evaluation result.

133. I am requesting clarification regarding the CMMC requirements from BAA HR001126S0010 Amendment 1 Section 3.3: for TA3:

1. For TA3, which Level 2 status satisfies the requirement: Level 2 (Self) or Level 2 (C3PAO)?

Level 2 (Self)

2. Does the Section 3.3 requirement also apply if the award type is an OT Agreement or Cooperative Agreement?

CMMC requirements will apply. Please note, in all cases, the Government Contracting Officer shall have sole discretion to select award instrument type, regardless of instrument type proposed, and to negotiate all instrument terms and conditions with selectees.

3. If all FCI and CUI will be processed, stored, and transmitted exclusively within a subcontractor-managed enclave covered by a current Level 2 CMMC UID, with prime personnel accessing it only through restricted VDI that prevents local storage, transfer, or printing, must the prime proposer's CAGE independently hold Level 2 status and a current affirmation in SPRS, or may the subcontractor's CMMC UID and status satisfy Section 3.3?

The (Prime) proposer must demonstrate compliance with the applicable CMMC level at the time of contract award and maintain that level throughout contract performance

134. The Volume I template (P1) lists "Technical Papers (Optional)" among the items excluded from the 25-page limit. However, the template itself contains no Technical Papers section, and Section 13 provides only "Bibliography (Optional) — [A brief bibliography, limited to 5 pages, may be provided with links to relevant papers, references, reports, etc.]" Accordingly:

1. Are optional technical papers permitted as part of a full proposal submission under this solicitation?

See Attachment P1 – DICE Volume 1 Template Instructions. Optional items may be added after Section 15.

2. If so, where are they uploaded, how many may be submitted, and what format or page limits apply?

See Attachment P1 – DICE Volume 1 Template Instructions.

3. Will their contents be evaluated as part of the proposal review, or are they provided for reference only?

Optional documents are used for reference.

135. If we are requesting a Procurement contract, do we need to submit the “CF Biographical Sketch” and “CF Other Support” forms?

According to Section 3.2.1 and Section 6, the "CF Biographical Sketch" and "CF Other Support" forms are required only for proposers requesting an OT Agreement or a Cooperative Agreement. They are not required for a Procurement Contract.

136. Is the “Official Transmittal Letter” that is listed in the Table of Contents of Attachment P1 included in the 25-page limit for the Technical and Management Proposal?

No

137. The DICE Volume I Proposal Template states in the instructions under section 7 Statement of Work: "Do not include more than two organizations per task/subtask." Is there a rationale for limiting the number of organizations on a task? We are proposing with a larger team, and some tasking would benefit from inputs from three organizations.

The rationale is to ensure clear lines of responsibility and accountability. While your team's workflow may be collaborative, for the formal SOW, you must designate a primary and at most one secondary organization responsible for delivering each task. You can describe the collaborative nature of the work within the technical volume, but the SOW must adhere to the two-organization limit per task to maintain contractual clarity.

138. The proposal template asks for the SOW with tasks broken down by Contractor Fiscal Year. We have usually written SOWs with tasks broken down by program phase, and then provided a cost break down by task per Contractor Fiscal Year. Can you either re-affirm that you really want the SOW with tasks per fiscal year, or clarify that SOW tasks may/should be written per program phase?

Follow the template's instruction to break down the SOW tasks by Contractor Fiscal Year.

139. What is meant by interaction steps in different contexts? Does this refer to interactions between agents within each team or between agents and the simulation? For example, p. 13 discusses the expectation of fewer interactions (messages) each phase ($O(N^2)$ to $O(n \log n)$ to $O(n)$), while Table 2 on p. 20 shows increased interaction steps each phase (5K, 50K, 1M).

Fewer Interactions ($O(N^2)$ to $O(N)$): This refers to the efficiency of coordination, specifically for adaptation. It measures the number of messages required to recover from a perturbation, relative to the number of agents (N). The goal is to move from an

DICE Q&A

inefficient protocol where everyone talks to everyone ($O(N^2)$) to a highly efficient one where communication is sparse ($O(N)$). Increased Interaction Steps (5K to 1M): This refers to the total scale of the mission. It is the total number of messages exchanged by all agents throughout the entire mission. This number is expected to increase dramatically as the number of agents and the complexity of the mission grow in each phase. In short, the program wants to see solutions that can handle a much larger total number of interactions (scale) while being much more efficient in how they use those interactions to adapt to change (efficiency).

140. Who is responsible for compute resources to run the TA1/2 agent adapter? For LLM-based agents, is TA3 or TA1/2 responsible for providing access to an LLM inference endpoint? For neural net-based TA1/2 agents running on GPU, is TA3 responsible for providing sufficient compute to run such agents?

TA3 is responsible for the compute resources to run the entire simulation, including the agents and the TA1/2 adaptors during the formal evaluation events.

141. Quantifying TA3 compute costs without foreknowledge of TA1/2 agent approaches and compute requirements is inherently uncertain. The variance could be an order of magnitude or more depending on TA1/2 dependency on the proportion of agents reliant on LLM, neural nets, MARL, variance in the simulation or TA1/2 solutions impacting the # of trials required, and other factors. How should TA3 proposals address this variance in their technical and cost volumes? Does DARPA have an expectation on the maximum budget for compute resources?

The BAA does not set a maximum budget. It acknowledges this uncertainty and wants TA3 proposers to address it by proposing an efficient and flexible architecture.

142. May we still submit an abstract for review before the August 25, 2026 full-proposal deadline, or is the abstract submission window strictly closed?

The abstract submission window is strictly closed.

143. For a TA1+TA2 procurement-contract path involving no CUI, we understand that a CMMC Level 1 self-assessment recorded in SPRS is the minimum pre-submission requirement under Section 3.3. Please confirm whether that satisfies the applicable pre-submission gate.

Yes, that is correct. Section 3.3 details the CMMC requirements. It states that "To be eligible for an award under this solicitation, the CAGE code associated with the proposal must have a valid, unexpired Final Level 1 for TA1/TA2 and Level 2 for TA3 CMMC status recorded in the Supplier Performance Risk System (SPRS)," and "Proposers must demonstrate compliance with the applicable CMMC level at the time of contract award and maintain that level throughout contract performance."

DICE Q&A

144. The Proposal Summary Slide template is not sized for 8.5x11 aspect ratio, so if it is printed to 8.5x11 paper, the content will be shrunk to fit. Is it acceptable to submit a Proposal Summary Slide in PPT format that is sized for 8.5x11 aspect ratio instead of the current aspect ratio (7.5x13.33) of the template file?

A. No, it is not acceptable.

145. Can you clarify whether "Schedule and Milestones" is included in the page limit as well?

"Schedule and Milestones" is included in the 25-page limit.

146. In the BAA amendment, language was added to ensure that those who target TA1/TA2 submissions, must be CMMC Level 1 compliant to be issued an award. However, it also notes that TA1/TA2 proposals will not be handling CUI. Can you please confirm if CMMC Level compliance is required if no CUI is involved?

Yes, CMMC Level 1 compliance is required even if no CUI is involved.