DARPA-PS-26-26
Virtual-Integrated Twin for Autonomous Lifesaving (VITAL)
Frequently Asked Questions
Version 1
10 April 2026

DARPA-PS-26-26 was published to SAM.gov on 9 April 2026

General Questions

1. **Q: Has the Program Solicitation (PS) been published? Where?**
   A: Yes, it was published on 9 April 2026 and can be found on SAM.gov.

2. **Q: Will the slides from the Proposers' Day be made available to the registered attendees or posted online?**
   A: Yes. Information relayed during the Proposers' Day will be made available on the VITAL program page at DARPA's website: https://www.darpa.mil/research/programs/virtual-integrated-twin-for-autonomous-lifesaving

3. **Q: Is teaming required?**
   A: While teaming is not required, it is strongly encouraged to provide the expertise and capabilities needed to achieve the VITAL program goals. Proposing teams should have a plan in place for managing team interactions and future technology transitions.

4. **Q: What is the expected award size range? How many awards are expected?**
   A: DARPA has not established predetermined award amounts. The number and value of awards will be contingent upon the quality and innovation of the proposals received.

5. **Q: What types of award instruments are anticipated?**
   A: Other Transaction for Prototype agreements.

6. **Q: Should abstracts address both Phase 1 and Phase 2 of the program?**
   A: Abstracts are required to detail the full technical team and present a complete research plan for both Phase 1 and Phase 2.

7. **Q: Do proposed approaches have to address both Focus Area 1 (FA1) and Focus Area 2 (FA2)?**
   A: Yes.

8. **Q: Can more than one (1) abstract be submitted from a single organization? Can more than one (1) abstract be submitted by the same researchers from a single organization?**
   A: Yes (to both scenarios), an individual/organization can be a co-investigator on multiple teams and must address both FA1 and FA2. If **invited to submit an oral proposal package**, a clear path will need to be established to ensure no conflicts are present between the efforts. Proposers who are on multiple teams should be cognizant of the distribution of the level of effort across multiple **proposals** .

9. **Q: Is there a preference for universities to propose as a prime versus a subcontractor?**
   A: No.

10. **Q: Are there any best practices, recommendations, or limitations for the size of proposing teams?**
    A: Proposers should choose the best team to achieve their proposed system.

Specific Questions

11. **Q: Is imaging a requirement for development of the high-fidelity models?**
    A: No. Imaging is not required. However, the construction of the high-fidelity (HF) models is a critical aspect of the work and must adhere to the following requirements:

    - Data Lineage: HF model development must be grounded in, and traceable to, relevant personalized medical data. The choice of data modality is left to the performer and should be driven by the medical condition and modeling approach. Potential data sources may include imaging, chemical or biochemical measurements, electrophysiology, physiologic waveforms, laboratory values, omics, clinical observations, or other relevant modalities. 'Image-to-Physics-to-Twin' was used as a shorthand for deriving mechanistic, patient-specific models from available medical data. Imaging could be one possible input. Performers should use whatever personalized data modalities are most appropriate for the condition and use case.

    - Scope: The HF models must accurately capture the full scope of the medical condition and intended use cases described by the performer, using whichever data modalities are necessary to represent the underlying physiology and pathology.

    - Intended Purpose: The HF models must be constructed with the explicit intent of enabling later reduction to a reduced-order model (ROM). Accordingly, performers should select data modalities and model representations that preserve the essential mechanisms, states, and parameters needed to support a credible and interpretable HF-to-ROM transition.

12. **Q: Are there any limitations on the types of sensing modalities and sensors that can be proposed (e.g., "wearables")?**
    A: No specific sensing modality is prohibited. Investigators are encouraged to consider and leverage whatever sensing modalities are most appropriate for the medical condition and modeling approach.

    However, this program is not intended to support the development of fundamentally new sensing modalities or advanced sensor hardware. Proposers should primarily rely on commercially available, off-the-shelf, or otherwise mature sensing technologies that can be acquired and integrated with minimal development effort. For example, if a performer determines that photoplethysmography (PPG) is required, they should plan to use an existing

PPG sensor rather than propose development of a novel ultra-low-profile, highly wearable PPG device.

Limited sensor adaptation or integration may be appropriate when necessary to support the measurement model, but the primary focus of the effort should remain on determining what measurements are needed, how they map to physiological state, and how they enable construction and reduction of the HF and ROM models—not on sensor invention or hardware maturation.

13. **Q: For Phase 1, can the High-Fidelity (HF) model focus on a specific chronic or acute condition selected by the proposing teams?**
APlease Section 1.4 of the PS. A single proposal may only address one condition, chronic or acute. Proposing teams may choose specific sub-categories for conditions; while they can still offer meaningful rigorous research plans for the objective of the program.

14. **Q: Is there interest in creating simulated patients (e.g., generate populations with specific types of conditions)?**
A: Performers may include simulated patients or limited synthetic populations if they determine that doing so is necessary to achieve the technical objectives of the project. The core premise of the program is to construct and evaluate models using data that are as personalized as possible to a single patient or biological instance. Accordingly, the primary emphasis is on developing, calibrating, and validating high-fidelity and reduced-order models from data associated with one individual instance rather than from large synthetic populations.

15. **A: Q: For Phase 1, can the High Fidelity (HF) model focus on a specific chronic or acute condition selected by the proposing teams?**

A: Yes. The solicitation identifies the broader chronic and acute condition categories of interest but proposing teams may select specific subcategories or representative conditions within those areas. Teams should choose conditions that enable a meaningful and rigorous research plan aligned with the program objectives, including construction, validation, and assessment of the HF model and its eventual transition to a reduced-order model (ROM).

16. **Q: Should proposers select one acute/chronic condition for Phase 1 and another for Phase 2?**
A: No. Per Section 1.4 of the PS, a single abstract may address only chronic **or** acute responses.

17. **Q: Should models be designed for both a healthy baseline and acute/chronic conditions?**
A: The proposer should make that decision.

18. **Q: Are acute versus chronic conditions more heavily weighted?**
A: No.

19. **Q: How important is explainability given that this is a "learn" program with no transition?**
    A: Please focus on metrics and establish a research plan in support of metrics.

20. **Q: Should the ROM be based on a single HF model and primitives or could there be some fusion across multiple HF twins?**
    A: The proposer should chooseto derive their technical approach.

21. **Q: The panel was talking about translation, including translation to edge cases; is that part of the program?**
    A: The panel discussion addressed potential future translation of digital twin technologies following demonstration of their capabilities. Performers should focus on establishing the scientific basis, performance limits, and credibility of the high-fidelity and reduced-order models. There is no requirement for the performers to focus on translation.

22. **Q: Are there any Technology Readiness Level (TRL) requirements for sensors at the beginning or end of the program?**
    A: No.

23. **Q: Is the ability for continuous updating applicable to the HF model or just the ROM?**
    A: Just the ROM.

24. **Q: How important is User Interaction/User Experience (UI/EX)?**
    A: UI/EX development/refinement is out of scope for this program.

25. **Q: How far into the future should the models be predicting? What is the expected prediction horizon?**
    A: Please see the metrics table in PS Section 1.5. The prediction horizon for hemorrhagic shock is $t \geq 15$ minutes. For heart failure Stage 4, $t \geq 80\% * (t_{onset\text{-}stage\ 4} - t_{onset\text{-}stage\ 3})$.

26. **Q: Should the ROM truly be operating in real time?**
    A: The appropriate level of "real-timeness" should be defined in the context of the metric being evaluated and the cadence at which new data arrive. For example, if the relevant sensing data are updated every 30 seconds and the required metric is a 5-minute early-warning prediction, then the ROM only needs to update fast enough to ingest the new data and produce a prediction within that window. Conversely, if the intended use is intervention guidance during a rapidly evolving hemorrhage or arrhythmia, the update rate may need to be on the order of seconds.

27. **Q: Does the program include transition/commercialization? Given that there's not a transition focus in the program, is there any advantage to including an industry partner on proposed teams?**
    A: No. Team composition is at the discretion of the proposer.

28. **Q: How should performers think about personalization of the models? Should the models be truly specific to an exact individual?**

A: Performers should view personalization as a means to achieve the required performance metrics, rather than as an end in itself. Models do not necessarily need to be fully individualized to an exact person if a less-specific representation is sufficient to meet the evaluation criteria.

The appropriate degree of personalization should therefore be driven by the metrics and use case. In all cases, the model should be personalized only to the extent necessary to achieve the required predictive accuracy, calibration, lead time, and intervention-response metrics.

29. **Q: What are the expectations by Phase for sensor development/fusion? When will multi-modal sensing be assessed? When should performers be using multi-modal data?**
A: Multi-modal sensing is not assessed as an independent program objective. Rather, sensing and sensor fusion are evaluated only in terms of how well they support construction, calibration, updating, and validation of the HF and ROM models. Performers should use whatever combination of modalities they determine is necessary at each phase to achieve their modeling objectives.

30. **Q: In Phase 1 are there any sensing modalities that are not allowed?**
A: There is no restriction for using any sensing modality for Phase 1, as long as they deliver on the project requirements for phase 2.

31. **Q: Is it necessary to include team members with expertise in medicine and advanced pathophysiology?**
A: The program does not prescribe specific team composition or required areas of expertise. However, proposers are expected to assemble a credible team with the expertise necessary to execute the proposed work. If the technical approach requires understanding of advanced medical conditions, physiology, or pathophysiology, then the team should include members with appropriate clinical or biomedical expertise relevant to those topics.