**DARPA-PS-26-26**
**Virtual-Integrated Twin for Autonomous Lifesaving (VITAL)**
Frequently Asked Questions

DARPA-PS-26-26 was published to SAM.gov on 9 April 2026
Batch 2 questions added on 16 April 2026
Batch 3 questions added on 21 April 2026
Batch 4 questions added on 23 April 2026
Batch 5 questions added on 29 April 2026

# BATCH 1 (4/10/26)

## *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?**
    **A:** Please 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.

# BATCH 2 (4/16/26)

1. **Q: Are performers going to be expected to touch on all 3 focus areas, or would a focus on a single area (image to model, 3D HF models, HF to ROM) be acceptable?**
   **A:** Abstracts must cover both Phase 1 and Phase 2, address both FA1 and FA2, and include both construction of HF models (FA1) and the HF-to-ROM / ROM construction effort (FA2). The PS explicitly indicates that a standalone ROM effort without a validated HF parent model is out of scope. Image-to-model is not a focus area.

2. **Q: Will DARPA provide data for model development? If so, will that be for both HF and ROM?**
   **A:** As indicated in the PS, DARPA will not provide data for model development. Proposers are responsible for the design of experiments and for acquiring all data needed to develop, personalize, and validate their HF and ROM models.

3. **Q: For Phase 1 do we need to cover both Heart Failure and hemorrhaging, or can we select one of them for Phase 1?**
   **A:** You select one condition, not both. A single abstract may address only chronic or acute responses. You should not switch one condition in Phase 1 and another in Phase 2.

4. **Q: What kind of human data would be acceptable for evaluation? Do we need to find and collect data from patients with chronic heart conditions or patients in emergency rooms with hemorrhagic conditions?**
   **A:** The PS emphasizes data lineage to relevant personalized medical data. The right dataset is the one that credibly supports the selected condition track and metrics.

5. **Q: For building the initial HF model, can we assume rich data from many sensors such as CT/ MRI etc. is available? Can data from perturbations be used? Or does the system need to build models from sparse and different subsets of sensors?**
   **A:** The Government will not be making data available for developing the models. However, performers are free to use rich data from many sensors. Imaging may be included but is not required per PS - Section 1.2 Program Description and Scope. Proposers should use whatever personalized modalities are most appropriate. Perturbation data are in scope. The key is that the HF model must be traceable to real personalized data and constructed with later ROM reduction in mind.

6. **Q: For ROM model, can we assume a HF model built from rich sensor data is available?**
   **A:** The ROM is expected to derive from an HF parent model. Whether the HF model uses rich or sparse sensing is up to the proposer, provided it supports the metrics and preserves the mechanisms needed for a credible reduction. The Government will not be making data available for developing the models. See Batch 2 FAQ, Q2 and Q5.

7. **Q: Do proposed approaches have to address both focus areas?**
   **A:** Please refer to the PS: each abstract must address both FA1 and FA2. At the same time, each abstract may address only one condition track -- either chronic or acute, not both. Thus, one condition track is permitted per abstract, but both focus areas must be addressed within that track.

8. **Q: Does the program expect or have preferences for each proposal/abstract to address the complete multi-physics HF model of the acute/chronic condition of interest (e.g., mechanics + hemodynamics + EP, etc), or is it within scope if a proposal only addresses a specific physiology (e.g., arrhythmia risk in chronic heart failure)?**
   **A:** The PS does not require "all physics" in the abstract. However, the HF model must capture the full scope of the medical condition and intended use case described by the proposer, and support rigorous construction, validation, and later HF-to-ROM reduction. A narrower subproblem can be in scope if it remains a meaningful, data-grounded, end-to-end VITAL effort rather than an isolated fragment.

9. **Q: Is the scope focused on Class III and Class IV?**
   **A:** The Phase 2 prediction metric is tied to forecasting Stage 4 onset relative to Stage 3 onset.

10. **Q: Is the intent to exclude acute decompensated heart failure but still include cases which may require hospitalization?**
    **A:** HF cases requiring hospitalization could be relevant if they are framed within the chronic HF track and associated metrics; what is out of scope is treating heart failure as the program's acute condition track.

11. **Q: Where can I find the abstract template and structure which is due on April 22?**
    **A:** The template and instructions are provided in the PS attachments. Attachments A and B are required for abstract submission, and abstracts must be submitted through BAAT.

12. **Q: Can the team be 3 persons: PI+2 professors/consultants? Is there any limitation to the number of team members?**
    **A:** The PS states that proposers should assemble the most appropriate team to achieve the goals and objectives of proposed effort and ensure its success. Team size and composition are otherwise at the proposer's discretion.

13. **Q: Can an individual be part of more than one team, and if yes, are there any constraints?**
    **A:** Yes. An individual or organization may be a co-investigator on multiple teams. When propose to be on multiple teams, proposer must demonstrate that there are sufficient resources dedicating to each project to ensure the success. If invited forward, there must be a clear path to mitigate conflicts of interest, and individuals participating on multiple teams should be mindful of effort distribution.

14. **Q: What specifically constitutes “AI techniques” in the context of Phase 2? Does this imply the construction of an AI model?**
    **A:** It is the responsibility of proposers to provide the most suitable solution to achieve the goals and objectives of VITAL program. Mechanistic AI is within scope. Other approaches are acceptable provided they support HF-to-ROM reduction and are not explicitly identified as out of scope in the PS.

15. **Q: Would an approach centered on cross-domain synchronization observability and physiological system stability, as a complement to digital twin prediction architectures, fall within the intended scope of the VITAL program?**
    **A:** It is the responsibility of proposers to provide the most suitable solution to achieve the goals and objectives of VITAL program and is not explicitly identified as out of scope in the PS.

DARPA-PS-26-26 Amendment 1 was published to SAM.gov on 21 April 2026

# BATCH 3 (4/21/26)

1. **Q: Do we need to include both animal and human data during phase 1? Or will the execution and validation of HF models based on large animal preclinical data be adequate for Phase 1?**
   A: There is no specific requirement to include both animal and human data. Performers should select the validation strategy that best supports program metrics and provides a sound and compelling approach for testing their proposed research.

2. **Q: Is REBOA considered within scope for a "perturbation"? Do we need to propose additional experiments for this mechanism? Is there a range of hemorrhage that is desired (e.g., 0-30% or 0-50%, controlled vs. uncontrolled?)**
   A: Performers are encouraged to select perturbations that are relevant to their proposed technical approach and program metrics.

3. **Q: Can we use "new/novel" sensors for the acquisition of data from the animal or human subjects? Or do we have to use current clinical tools only?**
   A: Teams are not limited to standard clinical tools. However, the program is not intended to fund development of new sensing modalities or new sensor hardware.

4. **Q: What is the allocated budget for this mechanism? In the solicitation, there is no explicit number, but there is mention of multiple awards will be funded for Phase 1. Is there a max. allowed budget for the proposal? Are indirect rates limited? Or allowed?**
   A: Multiple awards are anticipated. Performers should request the resources necessary to complete the proposed work and meet the program metrics. Cost realism is part of the evaluation criteria.

5. **Q: Is there a requirement to partner with a commercial entity for the transition of the work from Phase 1 to Phase 2 where real-time ROMs will be developed and tested?**
   A: No commercial or transition partner is required.

6. **Q: Can we include military research labs as partners for this proposal (e.g., ISR in San Antonio or USU/Walter Reed)? Or will they be used for external validation of the models developed?**

A: DARPA does not prohibit partnerships with military research laboratories. Performers should confirm independently that any proposed government entity is eligible to participate under the solicitation.

7. **Q: Are multiple PIs allowed? Can you share the estimated start date for this program?**
   A: Each proposal is expected to identify a single PI. The anticipated program start date is approximately November 2026 to January 2027.

8. **Q: Is the use of existing COTS equipment in a fundamentally new way considered a "mature sensing technology" (and therefore acceptable), or is it considered a "fundamentally new sensing modality" and thus not allowed?**
   A: Yes. Use of existing commercial or fielded devices, such as a cell phone or tactical radio, is acceptable provided it supports the program objectives and does not require development of a new sensing modality or hardware platform.

9. **Q: How do you define "mature sensing technology"?**
   A: There is no specific TRL requirement. However, the program will not fund development of new sensing modalities or new sensor hardware.

10. **Q: What does "novel imaging techniques" include/exclude?**
    A: New algorithms for interpreting or processing data from existing sensors are not excluded. However, performers should not propose development or validation of new imaging modalities or sensing hardware. The program is focused on construction and validation of computational models.

11. **Q: Are purely data-driven or "black box" machine learning models acceptable if being used to augment a baseline project but cannot be used to substitute for a baseline project?**
    A: Purely data-driven machine-learning models that rely only on statistical relationships between inputs and outputs, without an embedded and verifiable physical or biological mechanism, are out of scope. Machine learning may be used to augment the proposed approach, but the end model should not rely solely on a black-box ML method.

12. **Q: Are animal models only permitted within Focus Area 2, or can they be incorporated into Focus Area 1?**
    A: Performers should select and justify the experimental protocols and models most appropriate for supporting their technical approach and program metrics.

13. **Q: Are proposers permitted to use existing datasets and/or prospectively collect observational data, provided the data include all imaging, physiological waveforms,**

**time-stamped perturbations, and any other data needed to run the HF models and measure predictive accuracy?**

A: Use of existing datasets is not prohibited, provided they support the proposed technical approach and program metrics.

14. **Q: Could you please clarify the intent of the 3rd paragraph on page 5: "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."**

    A: The intent is to encourage patient- or instance-specific computational models rather than population-level simulators. Use of broader simulators is not prohibited, provided there is a compelling linkage to patient- or instance- specific data.

15. **Q: What is meant by 'large synthetic population data'?**

    A: Large synthetic population data refers to data generated from population-level simulators rather than from an individual-specific model grounded in real data.

16. **Q: Is output data from a real-world individual's HF model considered synthetic data?**

    A: Yes. However, it is individualized synthetic output grounded in and traceable to real personalized data, rather than being population level data. The primary concern is use of large population-based synthetic datasets.

# BATCH 4 (4/23/26)

1. **Q: Can DARPA clarify the relationship of non-imagining data to the three requirements related to translation of segmented anatomy that follow the statement in Section 1.2 about imagery?**
   A: Performers are encouraged to select data types that are relevant to their proposed technical approach and program metrics.

2. **Q: Are design and development of ROM-related technology permitted during Phase I (i.e., invoiced labor associated with a SOW task) or must all ROM activities only take place in Phase II?**
   A: Phase 1 efforts should focus on the Phase 1 metrics and deliverables. The HF model should be designed with eventual reduction to a ROM in mind, but the primary effort in Phase 1 should remain on HF model development, validation, and satisfaction of the Phase 1 metrics.

3. **Q: The Industry Day deck lists anything other than hemorrhagic shock (acute) and heart failure (chronic) as out of scope. Is that the definitive scope, or are those illustrative examples of acute and chronic cardiovascular conditions?**
   A: The only conditions of interest for this program are hemorrhagic shock (acute) and heart failure (chronic).

4. **Q: For Phase 1 validation, to what extent are performers expected to collect new experimental (human or animal) perturbation data versus leveraging existing datasets? Additionally, for Phase 2, are performers expected to conduct new intervention-focused experiments, or is it acceptable to validate intervention response primarily through simulation and/or retrospective data?**
   A: Use of existing datasets is not prohibited. The main requirement is that the data supports the proposed technical approach, model construction, calibration, validation, and program metrics. Proposers should present a compelling validation plan with credible data lineage and perturbation relevance. If existing or retrospective data are sufficient to support those goals, they may be used; if not, performers should propose additional data collection.

5. **Q: For the hemorrhagic shock condition, is there a desired/required amount of time AFTER shock occurs that the models should be evaluated?**
   A: Proposers should propose evaluation windows that are appropriate to the selected use case, available data, intervention timing, and program metrics. The validation plan should make clear which time horizons are being modeled and how prediction accuracy, directionality, magnitude, and latency will be assessed over those horizons.

6. **Q: Could you help clarify the third bullet point on page 7? Is it describing a requirement to continuously, automatically update model parameters to fit an individual's response by monitoring the individual's physiological state (thereby ensuring the models generate accurate personalized predictions)? What is the meaning of "inverse models" in this context? Is the inverse model intended to report the existence of a hemorrhage based on vital signs?**
   A: "Inverse models" are intended to determine what sensing is required to estimate the underlying physiological state and personalize the model to the individual. The underlying physiology should drive the sensing requirements and measurement models for personalization.

7. **Q: What model outputs will be assessed for high fidelity model performance for the hemorrhage use case?**
   A: Proposers should select outputs that are mechanistically relevant to hemorrhagic shock or the specific physiology of interest, supported by their data, and directly tied to the proposed validation plan and program metrics.

# BATCH 5 (4/29/2026)

1. **Q: Does the program expect a platform that enables the development of solution, or the solution itself or both?**
   **A:** DARPA is not asking for a generic platform in isolation. The expectation is a capability demonstrated through a concrete end-to-end solution (i.e., a working physiological twin for a defined use case) that also establishes a generalizable framework (HF to ROM to measurements to intervention prediction). In practice, this means:
   - Deliver a fully instantiated solution for a representative condition
   - While demonstrating that the underlying approach generalizes beyond that instance

2. **Q: How is the 'refined HF baseline' defined when computing the allowed ≤15% HF-to-ROM fidelity degradation in Phase 2?**
   **A:** The refined HF baseline serves as the reference model against which Phase 2 ROM performance degradation (≤15%) is evaluated for relevant and feasible scenarios. While some loss in fidelity is expected with any model order reduction, proposers should design their reduction strategy to minimize degradation while preserving the key physiological behaviors and predictive utility required to meet program metrics.

3. **Q: How will the ≥1000× runtime improvement metric be measured (reference hardware, solver configuration, wall-clock vs simulated time)?**
   **A:** The runtime improvement is expected to be measured using simulated time for equivalent simulation tasks, comparing high-fidelity (HF) and reduced-order model (ROM) implementations under well-defined scenarios. While DARPA does not prescribe specific hardware, proposers should define the reference hardware and solver configuration used in their evaluations and ensure an apples-to-apples comparison, with consistent fidelity targets across HF and ROM cases.

4. **Q: If the primary proposed condition remains chronic heart failure / structural heart disease, may a secondary validation vehicle for mechanism-preserving multiscale coupling, or would that be considered out of scope?**
   **A:** Yes, a secondary validation vehicle is acceptable if it directly strengthens or demonstrates the core mechanism. However, it should not dilute the primary use case and must be clearly positioned as supporting evidence rather than as a separate track.

5. **Q: Does DARPA expect all Phase 2 objectives and metrics (continuous updating, prediction horizon, intervention response prediction, expert concordance, etc.) to be demonstrated within one end-state twin, or can different metrics be evidenced across complementary demonstrators inside the same proposal?**
   **A:** Yes, a cohort of complementary models is acceptable; however, they should be unified within a single framework with a clear and consistent linkage to the HF models. The final deliverable should represent an integrated, end-to-end solution, rather than a collection of independent components requiring selection or assembly by the end user.

6. **Q: For Phase 1 perturbation validation in the chronic track, what classes of perturbations are acceptable for structural heart disease and device-planning use cases: must they be bedside physiologic perturbations, or can procedure-related, benchtop, ex vivo, or replay perturbations be used if they preserve causal response structure?**
   **A:** Proposers should select perturbations that are appropriate and clinically relevant to heart failure, and that support achievement of the program metrics.

7. **Q: Can the continuous-update requirement be satisfied using irregular or asynchronous clinical updates supplemented by limited sensor streams, or does DARPA expect evaluation primarily under truly continuous monitoring scenarios?**
   **A:** Yes, the continuous-update requirement can be satisfied using irregular or asynchronous updates, supplemented by available sensor data, provided there is a clear and compelling rationale that these measurements could ultimately be realized using sensors outside clinical settings.

8. **Q: What uncertainty quantification methods are acceptable for computing Expected Calibration Error (ECE), and how will ECE be evaluated?**
   **A:** DARPA does not mandate a specific uncertainty quantification method. Proposers may use any statistically sound approach, provided it produces well-calibrated uncertainty. ECE should be computed using standard methods and evaluated on held-out or blinded data.

9. **Q: For the heart-failure prediction-horizon metric, how will onset of Stage 3 and Stage 4 be operationally labeled, and may proposers define quantitative staging surrogates tailored to their selected chronic sub-condition?**
   **A:** Proposers may define operational criteria for Stage 3 and Stage 4 heart failure using clinically grounded frameworks or quantitative surrogates.

10. **Q: Will the Government-provided benchmark cases used for Phase 1 down-select be aligned to each performer's chosen chronic sub-condition and perturbation paradigm, or should proposers expect a common benchmark that may differ from their exact clinical demonstrated scenario?**
    **A:** Proposers should expect that Government-provided benchmark cases will be standardized to enable cross-performer evaluation.

11. **Q: How will '≥90% expert concordance' be measured?**
    **A:** The ≥90% expert concordance metric reflects agreement between model outputs and clinically relevant expert assessments. Proposers should define a clear evaluation framework and quantify agreement using appropriate metrics aligned with program goals, focusing on the most clinically meaningful decisions.

12. **Q: What constitutes '≥10× richer expert reasoning'?**
    **A:** The "≥10× richer expert reasoning" metric refers to the model's ability to provide more detailed, structured, and clinically meaningful insight than standard assessments.

13. **Q: At Month 36, when prototype models with associated code and data are delivered,**

- **what level of executability is expected if the solution depends on pre-existing commercial solvers and workflow components?**
- **is delivery of models, training data, wrappers, and interfaces sufficient, or**
- **is broader runnable capability on Government systems expected?**
- **how should proposers frame IP and data-rights assertions for pre-existing commercial software and libraries that are essential to the HF-to-ROM pipeline?**

**A:** At Month 36, DARPA expects delivery of an executable capability, including the models, associated code, and the necessary interfaces or wrappers to run the system. If the solution depends on commercial solvers or external components, proposers should clearly document all dependencies, execution requirements, and any licensing constraints. IP and data-rights assertions should distinguish between background IP and project-developed IP and describe how essential components will be made accessible to enable Government evaluation. These matters can be addressed in detail if the abstract is invited for a full OPP submission.

14. **Q: For a Hemorrhagic Shock focus, would DARPA accept human-subject data collected under an already-approved IRB protocol as a sufficient regulatory path to satisfy Milestone 2 (Month-4 regulatory approval), or is a program-specific IRB submission required?**
**A:** Yes, use of human-subject data collected under an existing IRB protocol is acceptable, provided the protocol can be amended as needed to meet DoD regulatory requirements, including submission to the appropriate DoD review bodies for any additional approvals.

15. **Q: Does DARPA consider a non-traditional defense contractor prime (under 10 U.S.C. § 4022) teamed with academic subawardees sufficient to satisfy the OT-P non-traditional contribution requirement, or must each phase independently demonstrate non-traditional contribution?**
**A:** Other Transactions for Prototypes must meet one of the following conditions:
    - At least one nontraditional defense contractor or nonprofit research institution is participating to a significant extent; or
    - All significant participants in the transaction other than the Federal Government are small businesses or nontraditional defense contractors; or
    - At least one third of the total cost of the prototype project is to be paid out of funds provided by sources other than the Federal Government.

16. **Q: Independent of the IP-assertion provisions in PS Section 6.5: if a performer intends to license existing IP, does the executed license need to be in place before abstract submission, before OPP submission, or before award negotiation?**
**A:** Before award negotiation.

17. **Q: The PS requires that a single proposal address both Focus Areas (FA1 — High-Fidelity models; FA2 — Reduced-Order Models). Phase 1 is predominantly HF-focused and Phase 2 introduces ROM. In the abstract, is a roughly 60/40 (FA1/FA2) page allocation appropriate given the phase structure, or does DARPA expect more balanced coverage?**
**A:** This is a determination for the proposer to make.