Nexus Concordat
Founder: Marjorie McCubbins
caelsereith@aetherprotocols.com
EIN 39-2932385 SAM.gov UEI WJ3SQQ24JHR4
July [DAY], 2026

Dockets Management Staff (HFA-305)
Food and Drug Administration
5630 Fishers Lane, Rm. 1061
Rockville, Maryland 20852

Re: Public Comment on Docket No. FDA-2026-N-4699 — Expedited Investigational New Drug Pilot Program; Request for Information (91 FR 37996, June 24, 2026)

Dear Dockets Management Staff,

Nexus Concordat respectfully submits this public comment in response to the Agency's Request for Information regarding the proposed Expedited Investigational New Drug (IND) Pilot Program. We thank the Agency for the opportunity to contribute to the design of a program whose success would materially advance the position of the United States as the preferred jurisdiction for first-in-human clinical research.

Statement of interest. Nexus Concordat is a Delaware C-Corporation (File 10605960, incorporated April 30, 2026, EIN 39-2932385) developing a biology-native compiled artificial intelligence substrate for use in regulated decision environments, including quantitative systems pharmacology (QSP) applications relevant to first-in-human dose selection. The architecture is the subject of four pending United States provisional patent applications: USPTO 63/939,190 (Substrate-Independent Memory Weighting Architecture, filed December 12, 2025); USPTO 63/962,385 (Neurochemical Language Model, filed January 17, 2026); USPTO 63/988,485 (AETHER: Adaptive Endocrine Transformer Heads, filed February 23, 2026); and USPTO 64/034,536 (Scyla: Biologically-Native Programming Language, filed April 9, 2026). A live specialist inference cell, `cardiac_gen3`, is publicly queryable at nexusneurodata.com/cell-demo.html with signed inference manifests verifiable by SHA-256 against a public lineage registry. Two prior comments have been submitted by Nexus Concordat in the same Operation TrialBlazer policy family: Docket

FDA-2026-D-6539 (*Quantitative Systems Pharmacology (QSP)-Based Dose Selection for MABEL in First-in-Human Trials; Draft Guidance for Industry*, filed June 25, 2026, comment tracking number `mqv-uiae-qwxn`) and Docket FDA-2026-N-4390 (*AI-Enabled Optimization of Early-Phase Clinical Trials Pilot Program; Request for Information*).

The Agency's proposal to establish a network of Qualified Research Institutions (QRIs) partnering with sponsors on Phase 1 IND development is materially relevant to the operating model contemplated by Nexus Concordat. We respond in the structure the RFI requests, identifying each question by its category and number.

Executive summary of positions. We offer the Agency two core recommendations, each developed in the section-by-section responses that follow.

1. **Establish a specialization tier for methodology-focused QRIs.** The eligibility criteria in Section I.C describe an integrated multi-disciplinary institution model. In practice, sponsors routinely assemble specialized providers rather than engage a single institution comprehensive across nonclinical, clinical, and CMC. A tiered qualification structure recognizing both *comprehensive* QRIs (meeting all four Section I.C criteria) and *specialized* QRIs (recognized for depth in one discipline, operating in partnership with complementary providers) would better reflect current practice. A methodology-axis specialization for quantitative systems pharmacology and model-informed drug development substrates directly supports the Quantitative Systems Pharmacology-Based Dose Selection draft guidance the Agency published under Operation TrialBlazer on the same date as this RFI.

2. **Preserve small-sponsor and specialized-QRI access through pilot design.** As written, the eligibility criteria structurally favor sponsors partnering with large integrated CROs. Small biotechnology sponsors and academic-spinout sponsors are unlikely to afford the fee structure of a comprehensive QRI. Including specialized QRIs at differentiated price points, and permitting partnership arrangements between specialized QRIs, widens sponsor access and increases the diversity of the pilot cohort.

Responses to Section II Questions

A. Pilot Program Design and Implementation

A.1. QRI Qualification and Capabilities

Response to A.1.i (Recommended changes to QRI capabilities, infrastructure, and leadership).
We recommend the Agency distinguish two QRI tiers under the pilot. *Comprehensive QRIs* meet all four capability criteria enumerated in Section I.C and provide end-to-end IND development support across nonclinical, clinical, and CMC. *Specialized QRIs* demonstrate depth in one of the three discipline areas or in an emerging cross-cutting methodology such as QSP, model-informed drug development (MIDD), model-informed first-in-human dose selection, or computational safety signal detection. Specialized QRIs would participate in the pilot in partnership with a comprehensive QRI, or with a complementary specialized-QRI ensemble, such that the sponsor engagement collectively covers all disciplines required for the IND submission even when no single institution does.

This bifurcation matches current industry practice. Sponsors frequently engage a full-service CRO for clinical operations and CMC while separately engaging a QSP or MIDD consultancy for computational input on dose selection. Requiring every pilot participant to be comprehensive across all disciplines would concentrate the pilot to a small number of large CROs and exclude methodology innovators whose value proposition is depth rather than breadth.

Response to A.1.ii (Additional required capabilities).
We recommend the Agency additionally consider the following capabilities where relevant to specialized QRI participation:

- Model-informed drug development substrate capability, including QSP modeling relevant to first-in-human dose selection under the Agency's contemporaneous QSP-MABEL draft guidance.
- Data integrity and reproducibility infrastructure, including 21 CFR Part 11 conformance for electronic records, ALCOA+ audit-chain integrity, and, for computational QRIs, reproducibility mechanisms such as signed inference manifests that permit byte-exact external verification of any recommendation of material influence to the IND submission.

- Interoperability capacity: demonstrated ability to consume outputs from one QRI as inputs to the sponsor's overall IND package, and to hand off intermediate products to complementary QRIs where the specialized model requires it.

Response to A.1.iii (Unfeasible or unattainable qualification considerations).
The criterion requiring a "track record of success in regulatory affairs, including supporting IND submissions across relevant therapeutic areas and product modalities" is challenging for institutions offering methodology innovations that did not previously exist during a full IND lifecycle. Novel methodology cannot demonstrate an IND track record without prior IND involvement, yet is excluded from IND involvement absent that track record. We recommend the Agency permit specialized QRIs to substitute the following in place of an IND-execution track record:

- Peer-reviewed publications describing the methodology and its validation.
- Prior Agency interactions in adjacent programs (Model-Informed Drug Development Paired Meeting Program, Biomarker Qualification Program, Innovative Science and Technology Approaches for New Drugs, prior public comment or workshop participation).
- Issued or pending patent applications documenting the methodology's originality.
- Public reproducibility artifacts, including signed inference manifests, open lineage registries, and cryptographic verification recipes, demonstrating that computational recommendations can be verified byte-exactly by the sponsor, the Agency, and third parties.

Response to A.1.iv (Different specializations by therapeutic area or modality).
We support therapeutic-area and modality specializations, and would extend the specialization dimension to include a third axis. We recommend three specialization axes, with QRIs permitted to qualify along any combination:

- Therapeutic area (oncology, central nervous system and psychiatry, cardiovascular, infectious disease, and so on).
- Modality (small molecule, biologic, cell or gene therapy, nucleic acid, combination product).

- Methodology capability (QSP and MIDD for first-in-human dose selection, computational safety signal detection, biomarker-informed enrichment, mechanism-of-action inference).

A methodology-axis specialization is directly aligned with the Agency's stated commitment under the New Approach Methodologies (NAMs) program and the April 2026 NAMs progress report cited in the Operation TrialBlazer announcement.

Response to A.1.v (Self-owned IRB and clinical trial site).
Specialized QRIs should not be required to own an Institutional Review Board or clinical trial site. Established partnerships with credentialed IRBs and trial site networks should be sufficient. Requiring in-house ownership limits specialized QRIs to large-CRO-adjacent providers and excludes methodology innovators, computational providers, and academic-affiliated methodology developers. For comprehensive QRIs the ownership requirement is more natural. For specialized QRIs it is a structural barrier without commensurate participant-safety benefit, because the specialized QRI is not itself conducting the clinical trial.

Response to A.1.vi (Dual-capacity QRIs serving as both IRB and regulatory advisor).
Dual-role QRIs should be required to disclose their concurrent roles to the Agency and to the sponsor at engagement onset, to adopt a documented information barrier between the IRB function and the regulatory-advisory function, and to permit Agency inspection of the barrier documentation as part of periodic QRI oversight. Reporting should include the QRI's written attestation that no cross-role influence has affected any specific recommendation.

A.2. Pre-IND and IND Review Process

Response to A.2.i (Recommended changes to QRI responsibilities and expected tasks).
For methodology-focused specialized QRIs, we recommend the responsibility set be scoped to the specific discipline the QRI supports rather than requiring participation across all three (nonclinical, clinical, CMC). A QRI providing QSP-informed dose selection input, for example, should be able to participate meaningfully at the nonclinical pharmacology and dose-justification layers without also being asked to review CMC content it has no comparative advantage in evaluating.

Response to A.2.iii (Independent, objective review while partnering with sponsors).
For computational and methodology-focused QRIs, we recommend the Agency consider

signed inference manifest emission as a supporting mechanism. When a computational QRI issues a recommendation derived from a model, the recommendation is accompanied by a cryptographic manifest identifying the model checkpoint, training corpus, sampling parameters, and input by hash. The Agency and any external auditor can re-execute the model with the manifest and verify the recommendation byte-exactly. This mechanism does not replace human independence; it supplies a reproducibility-based check that no comparable manual review offers.

A.3. Drug Eligibility and Participation

Response to A.3.i (Product classes best suited for the pilot).
We recommend the pilot include drug classes for which first-in-human dose-selection uncertainty is greatest and the Agency's contemporaneous QSP-MABEL draft guidance most naturally applies:

- Central nervous system drugs with novel mechanisms of action, where receptor-level pharmacology is well characterized but system-level dynamics are not fully predictable from preclinical data.
- Autonomic and cardiovascular agents where inter-species dose scaling is unreliable.
- Immunomodulators with subtle pharmacokinetic-pharmacodynamic relationships.

These classes are precisely where computational substrate provides the largest marginal value and where a QSP-methodology specialized QRI can most meaningfully contribute to a sponsor's IND package.

Response to A.3.ii (Prioritization considerations).
Prioritization should consider (i) alignment between the drug's mechanism and available QSP or MIDD methodology, (ii) severity of the target disease and adequacy of existing therapies, (iii) sponsor size, with small sponsors receiving priority to test whether the pilot expands access to the first-in-human pathway rather than concentrating it, and (iv) diversity of the resulting cohort across modalities.

Therapeutic areas of particularly high public-health priority in our view include combat-veteran post-traumatic stress disorder, treatment-resistant depression, and traumatic brain injury sequelae. Each combines high patient need, chronic under-development of new

therapies, and neurochemistry that computational substrate is uniquely positioned to characterize.

A.4. Pilot Scope and Scale

Response to A.4.i (Number of QRIs, therapeutic areas, and modalities).
We recommend an initial pilot cohort of at minimum three to five comprehensive QRIs (representing the traditional multi-disciplinary institution model) alongside three to five specialized QRIs distributed across the three specialization axes above, with at least one specialized QRI on the methodology axis so that the pilot can test whether QSP or MIDD substrate integration improves IND submission quality. Five to ten sponsor-QRI engagements across at least three therapeutic areas and at least two modalities would provide sufficient scale to detect variance in QRI performance while remaining small enough for the Agency to monitor rigorously.

A.5. Implementation and Feasibility

Response to A.5.i (Resources QRIs need to develop and maintain capabilities).
Resource requirements differ meaningfully by QRI type. For a specialized computational QRI operating on the methodology axis, typical parameters include:

- Staff: 3 to 15 full-time equivalents, including regulatory affairs, computational modeling, and quality and documentation personnel.
- Quality systems: alignment with IEC 62304 (medical device software lifecycle, applied to computational QRIs by analogy), ISO 13485 (quality management systems) where applicable, 21 CFR Part 11 (electronic records), and ALCOA+ data integrity principles.
- Technology infrastructure: computational infrastructure for model training and inference, cryptographic infrastructure for signed manifest emission and public lineage registry maintenance, and secure electronic records infrastructure.
- Estimated annual operating cost: approximately $1 million to $10 million depending on scope, roughly an order of magnitude below the annual operating budget of a comprehensive full-service CRO.

Recognizing the specialized-QRI cost profile explicitly matters for pilot design because

small sponsors are unlikely to engage a comprehensive QRI at typical full-service CRO rates but can meaningfully engage specialized QRIs at these more accessible price points.

Response to A.5.ii (Existing models and international examples).
Four examples are directly relevant to pilot design:

- The Center for Drug Evaluation and Research **Model-Informed Drug Development (MIDD) Paired Meeting Program**, expanded under PDUFA VII, already recognizes model-informed computational input into first-in-human decision-making. The proposed pilot could build on the MIDD program's demonstrated Agency-industry interaction patterns.
- The **Biomarker Qualification Program (BQP)** and the **Innovative Science and Technology Approaches for New Drugs (ISTAND)** program have established precedent for qualification of specialized methodology tools distinct from full-CRO qualification.
- The **New Approach Methodologies (NAMs)** work stream, with the April 2026 progress report demonstrating Agency commitment to computational and AI-powered approaches as a category.
- The **European Medicines Agency Innovation Task Force** process, which includes specialized methodology providers alongside sponsors in scientific advice proceedings, offers a comparable international precedent for tiered qualification.

B. Risks, Oversight, and Evaluation

B.1. Risks and Safeguards

Response to B.1.i and B.1.ii (Risks and safeguards).
Three risks warrant explicit consideration:

Over-reliance on QRI recommendations. A sponsor could treat a QRI recommendation as a substitute for its own scientific judgment rather than as expert advisory input. Mitigation: the Agency retains sole regulatory decision-making authority (as stated in Section I of the notice), and the QRI engagement agreement should explicitly state that QRI recommendations are advisory and do not diminish sponsor responsibility for IND content.

Methodology opacity. For computational QRIs specifically, sponsors and the Agency cannot always inspect why a model recommends a particular dose or safety margin. Mitigation:

require signed inference manifest emission for any computational QRI recommendation of material influence to IND content. The manifest permits the Agency and the sponsor to reproduce the recommendation independently and verify it byte-exactly. It also permits third-party audit without exposing model weights, corpus content, or other commercially sensitive material.

QRI-sponsor lock-in. If QRI-sponsor engagements become long-term and confidential by default, the pilot may inadvertently reduce sponsor mobility across QRIs. Mitigation: ensure QRI engagement agreements do not include exclusivity provisions or long-tail exit terms that limit sponsor mobility across the pilot cohort.

Response to B.1.iii (Inequitable access favoring well-resourced sponsors).
Yes, absent design mitigations. The eligibility criteria as currently written structurally favor large sponsors partnering with large CROs, because small sponsors are unlikely to afford comprehensive-QRI fee structures. We recommend the following mitigations:

- Include specialized QRIs in the pilot at differentiated price points, as discussed under A.1.i and A.5.i above.
- Permit partnership arrangements between specialized QRIs, so that a sponsor may engage two or three specialized QRIs collectively covering the required disciplines rather than a single comprehensive QRI.
- Consider reserving a portion of the pilot cohort for small-sponsor engagements, for example sponsors with fewer than 100 full-time equivalents or fewer than $50 million in annual research spend.

B.2. Oversight and Accountability

Response to B.2.iii (Transparency alongside sponsor confidentiality).
For computational QRIs, signed inference manifest emission offers a mechanism that resolves the tension between transparency and confidentiality. The manifest publishes cryptographic hashes of the model checkpoint, training corpus, sampling parameters, and input, sufficient for external verification. It does not publish the model weights, the training data, or the underlying corpus. The Agency and sponsors can reproduce and verify computational recommendations without exposing commercially sensitive information. Nexus Concordat operates such a signed-manifest infrastructure on the live cardiac specialist cell publicly available at nexusneurodata.com/cell-demo.html, and would

be pleased to demonstrate the mechanism to the Agency.

B.3. Success Metrics and Evaluation

Response to B.3.i and B.3.ii (Data collection and success metrics).
Recommended metrics include:

- *Time metrics*: time from sponsor QRI engagement to first submission component; time from full IND submission to safe-to-proceed notification; time from safe-to-proceed to first-in-human first-patient-dosed.
- *Quality metrics*: rate of clinical hold imposition; rate of Agency information requests during rolling review; sponsor-reported rework required post-review.
- *Safety metrics*: rate of dose-modification protocol amendments during Phase 1; rate of serious adverse events attributable to dose selection.
- *QRI performance metrics*: sponsor satisfaction, Agency reviewer satisfaction, adherence to QRI engagement agreement scope, on-time delivery of QRI written recommendations against the rolling submission cadence.
- *Access metrics*: distribution of pilot participants across sponsor size, therapeutic area, and modality; time from RFI publication to first small-sponsor pilot engagement.

Success should be defined as measurable improvement against pre-pilot baselines on time-to-first-in-human and clinical-hold rate, together with demonstration that the pilot expanded first-in-human pathway access rather than concentrated it.

C. Additional Considerations

Response to C.1 and C.2 (Additional or alternative approaches).
Complementary approaches the Agency could pursue in parallel with the pilot include:

- **Expansion of the MIDD Paired Meeting Program** to include specialized computational QRIs as recognized participants, permitting sponsors to bring specialized methodology input into pre-IND meetings without the pilot's full engagement structure.

- **Broadening of Pre-IND Meeting formats** to permit rolling submission of pre-IND briefing content, giving sponsors a lighter-weight iteration option than the pilot's structured QRI engagement.

The proposed pilot offers structural advantages over these alternatives when the sponsor benefits from sustained multi-discipline QRI collaboration across nonclinical, clinical, and CMC. It offers less advantage where one-off methodology input under an expanded MIDD process would be sufficient. Operating the pilot alongside expanded MIDD access maximizes the range of sponsors served.

Response to C.3 (Risk-based consideration of submitted information).
Product-risk stratification is appropriate. Higher-risk products (novel modalities, first-in-class mechanisms, drugs targeting rare and severe conditions) benefit most from QRI-augmented review because both sponsor uncertainty and Agency reviewer scrutiny are highest. Lower-risk products (for example, follow-on small molecules in well-characterized pharmacological classes) may derive smaller marginal benefit. Risk stratification could inform pilot prioritization but should not exclude lower-risk products where the sponsor meaningfully benefits from QRI support.

Response to C.4 (Considerations not addressed in the questions above).
Two additional considerations warrant Agency attention:

Data integrity infrastructure alignment. For computational QRIs, alignment with 21 CFR Part 11 (electronic records), ALCOA+ data integrity principles, and reproducibility mechanisms is central to trustworthy participation. We recommend the Agency recognize these capabilities as QRI qualification elements alongside the four already enumerated in Section I.C.

Open lineage registries. A public cryptographic registry of QRI computational artifacts (model versions, methodology iterations, corpus manifests) permits external verification, third-party replication studies, and post-pilot evaluation without exposing sensitive content. Nexus Concordat operates such a registry for the cardiac specialist cell at `patent.nexusconcordat.com/registry/`, and would be pleased to describe the operational model to the Agency if that would inform pilot design.

Closing

Nexus Concordat supports the proposed Expedited IND Pilot Program and thanks the Agency for the opportunity to comment. We would be pleased to demonstrate the specialized-QRI operating model described above, share technical details on signed inference manifest emission and public lineage registry infrastructure, and respond to any Agency follow-up questions the Dockets Management Staff considers helpful.

Respectfully submitted,

Marjorie McCubbins, BSc
Founder and Chief Executive Officer
Nexus Concordat
caelsereith@aetherprotocols.com