


The Sponsor Consensus

How Trial Sponsors Are Thinking About Real-Time Clinical Trials: An Industry Analysis of the Sponsor Response to FDA's RTCT Request for Information



Executive Summary

Based on the 20 public comments filed by drug and biologic developers to FDA's Request for Information on the AI-Enabled Optimization of Early-Phase Clinical Trials Pilot Program, the operational vehicle of the Real-Time Clinical Trials (RTCT) initiative. Positions are analyzed in aggregate and are not attributed to individual companies.

On April 29, 2026, FDA opened a public docket seeking input on a pilot program that would move clinical trial oversight from retrospective, episodic submission toward continuous, real-time review. By the time the comment window closed on June 29, the docket held 239 submissions. Inside that volume sits a smaller, more consequential dataset: the responses filed by drug and biologic developers themselves. Twenty sponsors, spanning the largest global pharmaceutical companies, mid-cap and specialty biopharma, emerging biotechs, and AI-native developers, put their thinking on the public record.

Paradigm Health analyzed all 20 sponsor responses in full, including attachments, and this report presents what the record shows in aggregate: where the industry converges, where emphasis genuinely differs, whether those differences cluster by organization size or therapeutic area, and what the whole picture signals about where early-phase development is headed. Four conclusions organize everything that follows.


First, the question has changed. Not one of the 20 sponsors argued against the initiative. Every response, including the most cautious, engaged with how to implement real-time trials responsibly rather than whether to pursue them. Several sponsors were already bought in to the point of spending their RFI response seeking absolute clarity on practical matters such as sourcing and validating a technology partner.

Second, trust has a specific architecture, and the sponsor community has effectively written its framework: source-data quality and provenance, human and patient oversight, a shared governance vocabulary, and enforceable confidentiality for models and data.

Third, the real divergences are about emphasis, and they cluster in revealing ways. Some cluster by organization size, some by therapeutic area, and one important divergence, the split between eager and cautious postures, does not cluster by size at all. It runs through the middle of large pharma itself.

Fourth, sponsors drew the map of where real-time pays and where it may not. The economics concentrate in the structural pauses between cohorts and phases. Several sponsors were candid that small, intensively monitored first-in-human studies may gain little. That honesty is itself a signal of how seriously the industry is planning.

The Record By the Numbers



Also in the record: 15 of 20 asked FDA to publish pilot learnings openly, 12 of 20 raised confidentiality or trade-secret protections, 11 of 20 prioritized AI-enabled recruitment and patient matching, and 9 of 20 named model drift monitoring explicitly.

Counts reflect positions explicitly raised in sponsors' filed comments, analyzed in aggregate. Source: public comments to Docket No. FDA-2026-N-4390, Regulations.gov, April 29 to June 29, 2026. Methodology on the final page.

Why This Docket Matters More Than Most: The Acceleration Race

The RTCT initiative is not an isolated experiment. It is one thread in a deliberate, multi-program push by FDA and HHS to compress the clinical development timeline, launched against an uncomfortable competitive backdrop. Public reporting through 2026 has documented that China now initiates more early-stage clinical trials than the United States, with FDA leadership stating earlier this year that China began roughly four times as many Phase 1 trials as the US in 2024, and HHS leadership writing that Chinese companies accounted for nearly half of global pharmaceutical licensing deals last year (remarks and commentary reported by FirstWord Pharma and The Hill, April and June 2026). In June 2026, HHS expanded its clinical development modernization push with Operation TrialBlazer, a program intended to shorten early

development timelines by six to twelve months through measures including rolling IND review, and framed explicitly as an effort to reverse the migration of Phase 1 trials overseas (The Hill, June 22, 2026).

The RTCT initiative, announced in April 2026, is the data-infrastructure arm of the same strategy: if trials are to move faster, the information flowing between sponsor and regulator has to move faster too, without losing its integrity. The 20 sponsor responses analyzed should be recognized as the industry’s first collective statement on what an accelerated, continuously supervised development model must look like to be trusted. That is why the patterns in this record deserve close reading: they will shape guidance, program design, and competitive positioning well beyond this single docket.

Who Showed Up: The Shape of the Sponsor Record

The 20 sponsor responses span the full structure of the industry, which is what gives the aggregate patterns their weight. Consistent with industry norms, this report characterizes positions at the segment level and does not attribute views to named companies; the underlying comments are public record. Therapeutic perspectives represented across the filings include oncology, rare disease and gene therapy, immunology, neurology, cardiology, metabolic disease, and RNA medicines.

Segment	Filings	Character of the responses
Large global pharma	6	The longest and most prescriptive filings, dense with governance mechanics, validation frameworks, and program-design recommendations.
Mid-cap and specialty biopharma	4	Focused and operational: precise questions about eligibility, data scope, signal reliability, and how to engage a technology partner.
Emerging and small biotech	5	Access-oriented: strong support paired with concern that participation may require infrastructure that only the largest sponsors possess.
AI-native developers	4	Frontier-oriented: proposals to expand the pilot’s scope, therapeutic reach, and technical ambition.
Academic-industry investigators	1	A systems view framing real-time trials as an evidence-infrastructure challenge as much as an AI challenge.

Where the Industry Converges

Five positions command majorities across every segment, and together they read as the industry's collective terms of engagement for a real-time model.

Support is unanimous, and it is practical. All 20 filings support the initiative. The tone is not aspirational but operational: sponsors asked how technology partners are validated and made transparent to FDA, how source-data quality is assured before a signal is acted on, how real-time signals coexist with complete-dataset expectations, and how the sponsor-to-regulator communication cadence works while a trial is live. When the questions in a public record read like diligence checklists, the organizations asking them are planning, budgeting, and evaluating partners.

High data quality is the precondition, not a feature. Fifteen of 20 filings conditioned real-time value on the integrity, completeness, or provenance of source data, and the argument arrived from every direction. The most engineering-minded submission proposed that every signal reaching FDA carry structured data-quality context: the completeness of the underlying data elements at the moment of signal generation, the latency from source capture to receipt, and any open queries on the contributing data, with a tracked correction rate to quantify how often signals are later retracted or materially revised. The most cautious submission warned that accelerating unprocessed data does not eliminate data-cleaning work, it relocates it downstream into live regulatory interactions. And a filing from a sponsor with direct proof-of-concept experience identified where the pain concentrates today, naming manual cleaning, validation, and review in electronic data capture systems as a large bottleneck of trial conduct.

Humans stay in the loop, and patients join it. Twelve filings insisted that AI support, not replace, qualified clinical, statistical, and regulatory judgment, with the firmest formulations ruling out autonomous dose-escalation or patient-management decisions entirely. Notably, the record

pushes past the familiar human-in-the-loop framing: one rare-disease filing proposed a patient-in-the-loop model in which patients participate directly in AI governance, grounded in the observation that rare-disease patients already face a diagnostic odyssey averaging five years before a trial is even a possibility. Another filing extended the logic to measurement, arguing that remote-monitoring use cases should be evaluated on patient burden and visit reduction, not sponsor-side timelines alone.

A shared governance vocabulary is forming. Half of the filings anchored expectations to the NIST AI Risk Management Framework by name, and six cited FDA's January 2025 draft guidance on AI in regulatory decision-making, urging the agency to operationalize its existing risk-based credibility framework rather than invent a parallel one. Pre-specification of analytical approaches, formal change control, drift monitoring, and fit-for-purpose validation recur across segments. The compliance conversation has already standardized around a small set of reference frameworks, which materially lowers the cost of alignment for sponsors who adopt them now.

Openness and protection, together. Fifteen filings asked FDA to publish and share pilot learnings openly, and twelve raised confidentiality, trade-secret, or intellectual-property protections. These are not in tension; the most sophisticated filings proposed mechanisms that deliver both, including tiered confidentiality in which methodological learnings are publishable while model architecture and product-specific results remain protected as trade secrets under existing regulation. The message to FDA and technology providers alike: transparency and IP protection must be engineered together, or participation will be limited to sponsors who control their full model stack.

Where Emphasis Diverges, and How the Differences Cluster

Disagreements in the record are not about direction. They are about emphasis, sequencing, and who bears the infrastructure burden, and they cluster in three distinct ways.

Cluster 1: By Organization Size

Size predicts what a sponsor worries about, with unusual consistency. The largest filers concentrated on governance machinery: confidential vendor-disclosure pathways modeled on the Drug Master File so third-party AI providers can share model documentation directly with FDA, defined pilot durations with milestone checkpoints, written technical feedback on fixed timelines, regulatory sandboxes for higher-risk uses, and commitments to translate pilot findings into draft guidance on a stated schedule.

The mid-cap cohort concentrated on entry mechanics: whether direct-from-source or sponsor-cleaned data is in scope, how data ownership and participant transparency work, how signal reliability is quantified, whether multi-region trials are eligible, and how a sponsor should select and validate a technology partner in the first place. These are the questions of organizations close enough to participation to need practical answers, without the internal machinery to answer them alone.

The smallest filers concentrated on access economics, asking FDA for a tiered participation structure in which sponsors with existing real-time infrastructure enter one way and scientifically well-suited sponsors without it enter through a shared platform with technology-partner support, paired with explicit conditions of platform portability and no single-vendor dependency. Tellingly, this ask did not stay at the small end: one of the largest filers independently proposed a supported-entry pathway for exactly this cohort, on the grounds that pilot learnings should not concentrate among the most AI-mature organizations. When the top and bottom of the market converge on the same structural request from opposite directions, program designers tend to listen.

Cluster 2: By Therapeutic Area

Where a sponsor works shapes what it wants the pilot to become. Rare-disease filings emphasized time economics and patient governance, arguing that in populations where diagnosis alone averages five years, compressing development pauses is a patient-welfare issue, not an efficiency issue. A cardiology-focused filing made the case that cardiovascular disease suits real-time inference unusually well, combining decades of outcome-trial priors with continuously measurable physiologic signals. A metabolic-focused filing supplied the sharpest arithmetic in the record: trials built around hard endpoints in metabolic disease face observation windows of three to five years, and real-time transmission of validated surrogate endpoints could compress the wait to twelve to eighteen months, a proportionally larger saving than in oncology, where designs already lean on six-to-twelve-month surrogates. Oncology-focused filings, for their part, pushed toward adaptive designs: resistance biology, combination regimens, and AI-derived biomarkers governing eligibility itself.

The pattern is consistent: every therapeutic constituency argued that the pilot should expand toward its own disease area, and each supplied a defensible scientific rationale. Twelve of 20 filings urged broad therapeutic scope overall. For FDA, that is pressure to generalize quickly. For sponsors, it is a preview of the second wave: the record suggests demand for real-time architectures is already forming well beyond the oncology proving ground.


Cluster 3: The Divergence That Does Not Follow Size

The most strategically interesting split in the record runs through the middle of large pharma itself. Among the six largest filers, some wrote what amount to draft operating manuals, proposing pilot durations, statistical acceptance criteria, and guidance timelines, while others devoted their filings to defending the scientific value of what real-time models might displace: data cleaning, adjudication, and the structured FDA meeting process, warning against new review gates and against acting on premature, noisy signals.

Both camps support the initiative. What separates them is a strategic judgment about timing: move first and shape the rules, or let the eager absorb the first-mover risk and adopt what survives.

This matters because it means readiness posture is a choice, not a consequence of scale. Two organizations with similar resources, similar pipelines, and similar regulatory sophistication have visibly chosen opposite clocks. Every sponsor reading this record should recognize that the same choice now sits in front of them, and that it will be read by partners, investors, and the agency as a signal of intent.

How the Differences Cluster



Where Sponsors Located the Value, and Where They Honestly Did Not

The Economics Live in the White Space

Strip away the technology discussion and the record contains a surprisingly unified theory of value: the largest returns come not from making any single activity faster but from collapsing the structural pauses between activities. One large multinational named the target directly, the white space of drug development, and enumerated where it concentrates: gated pauses between dosing cohorts, the stop-start transition from adult to pediatric development, movement of oncology drugs into earlier lines of therapy, and, looking past the pilot's early-phase scope, converting accelerated approvals to traditional approval the moment an event-driven confirmatory trial crosses its threshold rather than months later. A smaller filer called the same phenomenon dead time and asked FDA to establish pre-pilot phase-transition benchmarks so reductions can actually be measured. The generalizable rule for any portfolio: the longer a program's structural dead time, the more a real-time architecture is worth to it. That equation, not technology enthusiasm, is where deployment decisions should start.

The Honest Counterpoint: Where Real-Time May Not Help

A credible plan needs a map of where not to start, and sponsors supplied one. Multiple filings, including from the largest segment, observed that small dose-escalation cohorts with limited data are already intensively monitored by humans and may see little efficiency gain from real-time submission, and could even be slowed if incomplete data is assessed without context. Others noted that first-in-human studies are hard to operationalize as continuous real-time trials precisely because no prior safety data on the agent exists and amendment likelihood is high, suggesting lower-risk, data-rich settings as better early proving grounds. Some would defer autonomous dose selection and AI-generated clinical recommendations entirely until validation foundations mature, and at least one would exclude generative AI for regulatory document drafting from the initial scope altogether.

These are not objections; every filer supports the initiative. They are the boundaries of the current value map, drawn by the organizations that will run the trials. The composite picture: real-time architectures earn their keep where data is dense, signals are continuous, and structural dead time is long, and they should be introduced carefully where datasets are sparse and human monitoring already saturates the problem.

What Sponsors Want FDA To Answer Next

Beyond stating positions, sponsors used the docket to put specific open questions on the agency’s desk. These asks preview the guidance and program-design decisions ahead, and several directly shape what participation will require of a sponsor organization.

The ask	Raised by	Why it matters to sponsors
Keep the pilot broadly applicable, not oncology-only	A majority (12 of 20), across all segments	Determines whether cardiovascular, metabolic, immunology, and rare-disease programs can enter early or must wait for a second wave.
Create a confidential third-party AI vendor pathway, modeled on the Drug Master File	Two of the largest filers, independently	Formalizes how technology partners plug into regulatory submissions without forcing IP disclosure through the sponsor.
Provide a tiered or supported entry route for sponsors without in-house infrastructure	Smaller developers, seconded independently by one of the largest filers	Decides whether participation is limited to organizations with bespoke pipelines or opened via shared, portable platforms.
Clarify multi-region eligibility and pursue international regulatory alignment	Multinational filers across size tiers	Most early-phase programs enroll across regions; a US-only model would not generalize to the trials sponsors actually run.
Clarify whether post-marketing and open-label extension studies are in scope	A mid-cap developer	Signals appetite to treat real-time transmission as lifecycle infrastructure, not an early-phase experiment.
Define how real-time signals coexist with complete-dataset expectations	Sponsors across segments	The operational crux: what a signal obligates a sponsor to do, and when, before the full cleaned dataset exists.
Publish examples of use cases FDA found insufficient	A large filer	Would give every sponsor a public map of where the validation and governance bar actually sits.

The International Dimension

Because clinical development is global, several filings treated international coherence as a make-or-break design constraint rather than a diplomatic nicety. Specifically, it was emphasized that without coordination through international regulatory bodies, sponsors could face contradictory obligations, one regulator demanding immediate data flow while another prohibits unplanned interim disclosures. A second filing grounded the same concern in engineering reality, noting that non-US electronic health record ecosystems differ materially in standards maturity, including

wide variance in FHIR adoption, and in consent frameworks, so a pilot that never demonstrates real-time signal flow across that heterogeneity cannot credibly claim to be a foundation for scaled implementation. One of the largest filers went further, asking FDA to make international sharing a named pilot deliverable with findings routed to ICH partners on a defined schedule. For multinational sponsors, the practical implication is to design real-time data architecture for regulatory plurality from day one.

The Measurement Vocabulary Is Forming Now

Perhaps the most quietly consequential content in the record is the set of metrics sponsors proposed, because whichever measures the pilot adopts will become the benchmarks every participant is judged against. A recognizable framework is already assembling across the filings.

Emerging measure	Where it appears in the record	What it captures
Decision concordance	Seven filings; the most detailed specifies Cohen's kappa for categorical decisions and intraclass correlation with confidence intervals for continuous ones, via parallel decision tracking	Whether AI-supported decisions match what traditional methods would have concluded on the same data, before anyone relies on them alone.
Model drift and change control	Nine filings name drift; some add notification triggers for retraining and vendor-initiated model updates	Whether a system that was valid at qualification is still valid mid-trial, and whether FDA hears about material changes.
Signal latency	Multiple filings; one urges publicly reported latency benchmarks across all pilot trials, measured from source capture to FDA receipt	How real the real-time actually is, measured end to end rather than claimed.
Post-signal correction rate	One mid-cap filing	How often transmitted signals are retracted or materially revised as data accrues: the single best test of whether speed is costing reliability.
Decision-quality outcomes	Several large and mid-cap filings	Whether faster decisions were also good decisions: dose selections confirmed downstream, reversal rates, and time from data availability to decision.
Patient-side burden	One mid-cap filing	Visit reduction and participant burden for remote monitoring, so efficiency is not measured on the sponsor side alone.

The strategic point is timing. Sponsors who begin instrumenting these measures internally will enter participation conversations with a solid foundation in hand and will be positioned to shape their participation strategy faster.

What the Record Suggests You Ask Your Own Organization

Distilled from the questions the 20 sponsors asked FDA, and from the gaps their comments reveal, six questions form a practical readiness check for any development organization watching the RTCT initiative:

- Could our trial data flow to a regulator in near real time today, and if not, which specific systems and processes are the constraint?
- If a real-time safety or efficacy signal fired tomorrow, who in our organization would see it, interpret it, and decide, and with what data-quality context in front of them?
- How would we validate and document a third-party technology partner to FDA's satisfaction?
- Which of our upcoming programs has the longest structural dead time, in cohort transitions, phase gaps, or endpoint observation windows, where compressed white space would be worth the most?
- Are we instrumenting the measures the pilot is likely to adopt, including concordance, latency, drift, and signal-correction rates, or would we be starting from zero?
- If our programs are multi-region, does our data architecture anticipate regulatory plurality, or does it assume a single regulator's expectations?

Organizations that can answer these six questions crisply resemble the rule-shapers in this record. Organizations that cannot are, for now, in the access-seeking cohort, and the record itself points to their fastest path forward: shared, standards-based, portable infrastructure rather than a multi-year internal build.

The “So What”: Reading the Record From the Inside

Paradigm Health is the technology partner selected for the RTCT initiative by the FDA and co-developer, with the agency, of the Clinical Trial Reporting Schema (CTRS), the standards-based layer that transmits trial signals from source to FDA. We have kept that perspective out of the analysis above deliberately; the patterns in the record speak for themselves. But readers deserve to know how the findings look from inside the initiative, so we close with four observations.

The validation questions already have an answer in production. The single most common thread in the record, how real-time signals are validated and how data-quality context travels with them, is precisely the problem CTRS was built to solve. Each signal transmitted under the schema carries its provenance and data-quality context, including the completeness and recency of the underlying data, so a real-time signal arrives with the information a reviewer needs to weigh it. The data-quality monitoring layer sponsors asked FDA to require is, in effect, a description of the transmission standard the initiative already runs on.

The access question is smaller than it looks. The record’s most poignant divide, between sponsors with bespoke pipelines and scientifically strong sponsors without them, assumes real-time capability must be built.

A shared, standards-based, portable transmission layer already exists, and it was designed so that participation is a matter of connection and enrollment rather than multi-year engineering. The tiered-access structure the smallest and largest filers converged on is, practically speaking, an adoption question, not an invention question.

The emerging measurement framework is native to the architecture. Latency measured from source capture, provenance carried with each signal, and correction tracking as data accrues are not aspirations to be bolted on later; they are properties of a well-designed transmission standard. Sponsors who want to enter the measurement conversation with evidence rather than intentions can begin instrumenting against these measures now.

In a race measured in months, architecture is the answer that keeps integrity. The competitive backdrop that produced this initiative rewards speed, but the sponsor record is unambiguous that speed without data integrity is not a prize worth winning. The way to compress development timelines without relocating risk downstream is infrastructure that makes fast and trustworthy the same property. That is the design brief the sponsor community has now written in public, and it is the brief this initiative was built against.

Ready to explore participation? Start with a Sponsor Briefing.

We will walk through the FDA pilot model, how your study portfolio fits, and what Paradigm Health’s partnership looks like from day one. Contact your Paradigm Health representative or visit paradigmhealth.ai.

Methodology

This report analyzes the 20 public comments filed by drug and biologic developers (trial sponsors) in response to FDA's Request for Information, AI-Enabled Optimization of Early-Phase Clinical Trials Pilot Program, Docket No. FDA-2026-N-4390, published April 29, 2026, with the comment period closing June 29, 2026. All 239 docket submissions were reviewed to isolate the sponsor subset; technology vendors, CROs, trade associations, advocacy organizations, and individual commenters were excluded from the counts in this report. Sponsor positions were coded from the full text of each filed comment, including attachments, as posted to Regulations.gov, and every quantified count

was verified against the source text. Counts (for example, 15 of 20) reflect positions explicitly raised in a sponsor's own filing; absence from a count means the sponsor did not raise the point, not that it disagrees. Consistent with industry norms, positions are aggregated and characterized at the segment level; this report does not attribute views to named companies, although all underlying comments are matters of public record and can be retrieved from Regulations.gov under the docket number above. Market-context statements regarding clinical trial volumes and FDA acceleration programs are drawn from public reporting as cited in the text.

This report is provided for general information and reflects Paradigm Health's analysis of public regulatory filings and public reporting; it does not constitute regulatory or legal advice.

