PHASE V
Adaptive. Causal. Intuitive.

PhaseV Biopharma Survey Report

The State of Clinical Trial Data Handoffs: Spend, Timelines, and Vendor Dynamics

June 2026 • N = 51 senior pharma and biotech leaders

~$3M	88%	76%
median external workflow spend, Phase I–III combined	of sponsors report during/post-trial workflows take 5+ weeks	of sponsors coordinate up to 3 vendors per trial for these workflows and 12% up to 5

Why This Matters

Clinical development runs on a sequence of handoffs, from protocol to SAP, SAP to SDTM/ADaM, data to TLFs, TLFs to CSR. Each is typically executed by a different internal team or external vendor. This survey set out to measure what that fragmentation actually costs sponsors in time and dollars. The findings, drawn from 51 senior leaders across clinical operations, biostatistics, statistical programming, and regulatory affairs, put a number on a problem most pharma and biotech teams feel but rarely quantify: reporting and documentation workflows are a multi-million-dollar, multi-month line item that scales with trial complexity, and does so predictably enough to model, and target, with automation.

Methodology

PhaseV commissioned an independent research firm to survey 51 senior leaders at top pharmaceutical and biotechnology companies in June 2026, spanning clinical operations, biostatistics, statistical programming, regulatory affairs, and clinical development. The survey covered pre-trial workflows (study synopsis, protocol, schedule of assessments, and CRF development) and during/post-trial workflows (SAP development, data review, SDTM/ADaM development, statistical programming, TLF generation, and CSR authoring), examining timelines, vendor utilization, operational bottlenecks, and external spend. Spend figures reflect a subset of 26 large biotech/pharma respondents. All findings are self-reported. PhaseV did not influence responses and has not independently validated the data.

Finding 1: Reporting timelines run long and get longer as trials mature

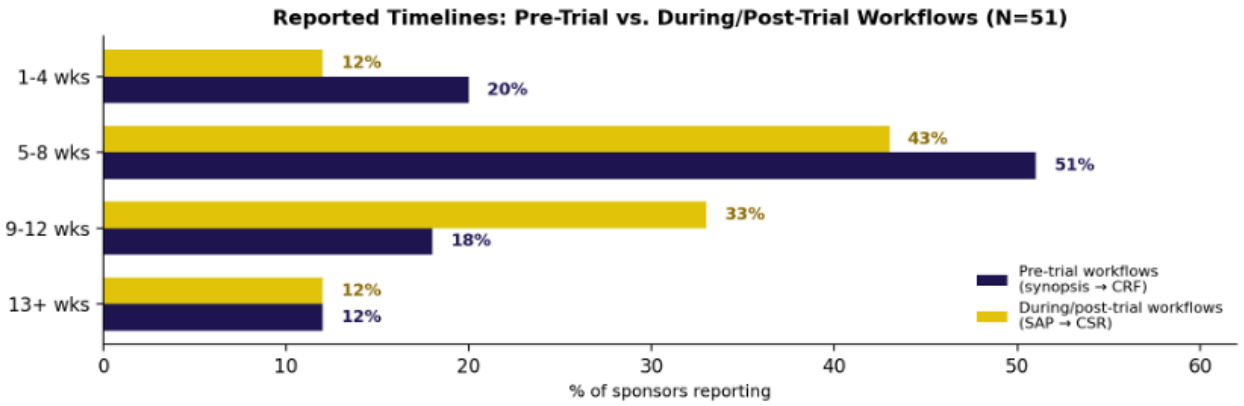

Figure 1. Sponsor-reported timelines for pre-trial vs. during/post-trial workflows (N=51).

Pre-trial workflows already run long: 51% of sponsors need five to eight weeks to move from study synopsis to CRF, and 30% need nine weeks or more. But the more consequential bottleneck sits downstream. During- and post-trial workflows (SAP through CSR) are where the real drag concentrates: 88% of sponsors report these critical-path activities take five weeks or longer, and 45% report nine weeks or longer. This is the stage where data volume, query resolution, and regulatory formatting requirements compound, and it is also the stage closest to database lock and submission, meaning delays here have the most direct line to trial cost and time-to-market.

Finding 2: Spend Scales With Complexity, Predictably Enough to Invest in

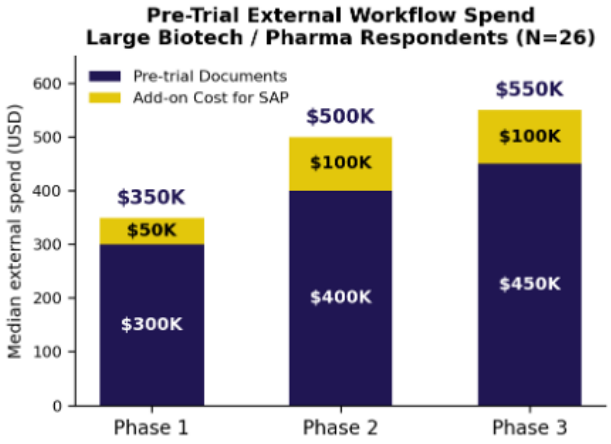

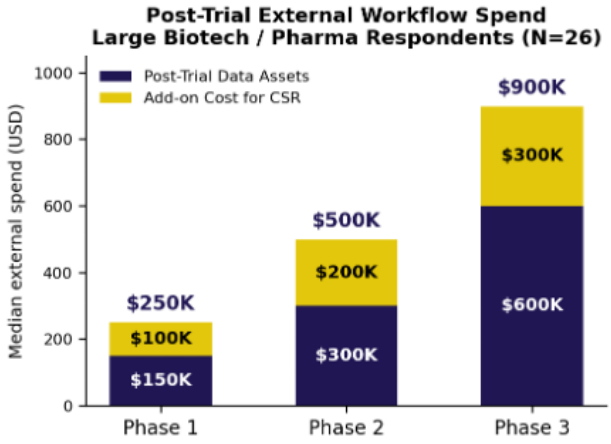

Figures 2–3. Median external workflow spend by phase, large biotech/pharma respondents (N=26).

Among large biotech and pharma respondents, median pre-trial spend rises from $350K in Phase I to $550K in Phase III, a 57% increase driven by greater protocol complexity, patient enrollment, and site expansion. Post-trial spend rises even more steeply, from $250K in Phase I to $900K in Phase III, a 260% increase driven by larger datasets, more complex statistical analyses, and expanded regulatory submission requirements. Combined across all three phases, sponsors report roughly $3M in external workflow spend over the life of a single program. The consistency of the pattern, spend scaling predictably with phase and data volume, is itself the finding that matters most for a buyer: this is a cost structure that can be modeled, benchmarked, and, PhaseV argues, standardized.

Finding 3: Fragmentation Is the Multiplier

Cost and timeline pressure alone would be a documentation-efficiency problem. What turns it into an operational risk is fragmentation: 76% of sponsors coordinate up to three vendors per trial across these workflows, and 12% manage four to five. Layer in an average of up to four protocol amendments per Phase II/III trial, each one triggering another round of manual updates across every downstream document and vendor, and the result is a process where delay is structural, not incidental. Vendor count and amendment frequency are the two variables sponsors have the least direct control over and, this data suggests, the two most correlated with timeline overrun.

What This Means for Buyers and Users of AI Conductor

For a VP of Clinical Operations, Head of Biostatistics, or Regulatory Affairs lead evaluating this category, three things should carry more weight than the topline dollar figure:

- Predictability is the real product, not just speed. The steep spend curve from Phase I to Phase III is not primarily caused by a single document produced too slowly, it is driven by rising volume of work and the growing number of vendor and stakeholder handoffs as trials mature. A platform's value should be

measured by how much of that curve it flattens, not just how many hours it saves on a single document.

- Vendor count is a leading indicator worth tracking internally before evaluating any vendor consolidation platform. Teams running 4-5 vendors per trial (12% of this sample) should expect proportionally higher integration and change-management costs from any new system, including AI Conductor, and should weight implementation support accordingly.
- Amendment-driven rework is where AI-native tools should show the clearest ROI. Because each protocol amendment cascades into downstream documents, a platform that can propagate a single validated change across artifacts, rather than requiring re-authoring at each step targets the compounding cost directly, not just the initial drafting time.

From an AI-capability standpoint, the distinction that matters for procurement and compliance review is between automated change propagation (auditable, low-risk) and generative document authoring (higher-value, but requiring human validation at each output). Buyers evaluating any platform in this category — including AI Conductor — should ask vendors to be explicit about which category each capability falls into, and what the human-review checkpoint looks like before a document reaches submission-ready status.

Conclusion

The survey does not prove that automation is the only fix for these bottlenecks, it establishes that the underlying cost structure (a $3M, multi-vendor, amendment-sensitive process) is large, measurable, and consistent enough across respondents to be a legitimate target for platform investment. That is the case this data makes, and the one buyers should hold any vendor in this category to.