Scaling Interactive Applications for Clinical & RWE Insights

Best Practices for Adoption, Compliance, and Long-Term Impact

Chris McSpiritt
Domino Data Lab
PhUSE 2026

Introduction

Interactive applications are changing how clinical and real-world evidence (RWE) teams deliver insights. Instead of waiting for static reports at predefined milestones, stakeholders can now explore results in real time—filtering by subgroup, adjusting time periods, and drilling into the data that matters most to their decisions.

Yet adoption of interactive apps across the life sciences industry remains inconsistent. Teams build one-off dashboards that gain traction with a handful of users, then struggle to scale the approach across therapeutic areas, geographies, or regulatory contexts. Without clear use cases, intuitive design, and the right governance framework, promising tools become shelfware.

This paper shares best practices for making interactive applications effective and sustainable across organizations. Drawing on lessons from clinical trial reporting and RWE dashboards, it outlines how teams can identify the right use cases, design for adoption across diverse stakeholders, and balance self-service interactivity with regulatory requirements.

The Insight Delivery Gap

RWE and clinical teams today face a convergence of pressures: exploding data sources, rising regulatory demands for evidence, and growing payer expectations for real-world outcomes data. Yet the tools most teams rely on are failing to keep pace.

What's Broken Today

- Static outputs dominate: stakeholders wait days or weeks for re-analysis when a new question arises, creating bottlenecks that slow decision-making at every stage.
- Fragmented tools across teams and studies mean that one group's Shiny app doesn't talk to another's SAS outputs, and institutional knowledge is lost between projects.
- Manual, undocumented workflows create compliance risk. When the path from raw data to a final figure exists only in someone's head, auditability becomes aspirational rather than real.
- There is no chain of custody from data to decision. Regulators, payers, and internal stakeholders increasingly demand traceability that current processes simply cannot provide.

What Great Looks Like

Organizations that get this right share a common pattern: they move from delivering reports to delivering applications. The difference is profound. A report answers the questions you anticipated. An application lets stakeholders answer the questions you didn't.

The target state combines interactive apps for real-time exploration, a unified environment that bridges RWE and clinical workstreams, reusable and governed workflows with built-in trust, and full traceability from raw data to insight.

When Do Apps Add Value Over Static Outputs?

Not every analysis needs an app. Building and maintaining interactive applications takes more effort than producing a static deliverable, so the investment must be justified by a clear return. A simple decision framework helps teams focus where apps drive the most impact.

Three Signals That an App Will Add Value

1. **Repeated questions.** Stakeholders repeatedly ask for the same analysis with different parameters—subgroups, time periods, endpoints. If the answer is "run it again with different filters," that's a signal to build an app.
2. **Multiple audiences.** Clinical operations, medical affairs, regulatory, and commercial teams all need the same underlying data but from different angles. An app serves them all without requiring separate outputs.
3. **Evolving data.** Data accumulates over time—interim analyses, ongoing RWE studies, post-market surveillance. An app that refreshes with new data eliminates the cycle of regenerating static reports.

Conversely, if an analysis is truly one-time, has a single well-defined audience, and the data is fixed, a static report may be perfectly adequate. The goal is not to replace all reports with apps, but to deploy apps where the return on investment is highest.

Designing Apps That Drive Adoption

Building an app is straightforward. Getting people to use it is the hard part. The most common failure mode is not a technical one—it's building an app that mirrors the developer's mental model rather than the stakeholder's decision process.

Start with the Decision, Not the Data

The single most important design principle is to orient the app around the decision stakeholders need to make. What action will this enable? A safety review board doesn't want to explore raw adverse event tables—they want to see whether a signal warrants further investigation. A medical affairs team doesn't want a general-purpose cohort builder—they want to understand treatment patterns in a specific population.

Progressive Disclosure

Show the essential view first. Let power users drill deeper. Don't overwhelm everyone with every parameter upfront. The best apps present a clean default state that immediately answers the most common question, then provide layers of interactivity for users who need more granularity.

Know Your Personas

A medical director, a biostatistician, and a regulatory reviewer have fundamentally different needs. An app that tries to serve all three with a single view will satisfy none. Persona-aware design—whether through role-based views, configurable dashboards, or progressive complexity—is essential for cross-functional adoption.

Consistency and Standards

Reusable templates, shared design systems, and consistent navigation patterns reduce the learning curve for every new app. When users encounter a familiar interface, they spend their cognitive energy on the data rather than figuring out how the tool works.

The Platform Shift: From Models to Applications

For years, the data science workflow in life sciences has ended at the model. A team would build an analysis, validate it, produce a report, and hand it off. The insights were locked inside code that only the original authors could interpret.

The emerging paradigm is fundamentally different. Instead of delivering models that produce insights, teams now deliver applications that enable decisions. This shift has three stages:

Stage	Activity	Description
1	**Build the Model**	Data scientists develop and validate analytical models in a governed, reproducible environment with full version control.
2	**Publish as an App**	Turn validated models into interactive applications—deployed as apps, APIs, or model endpoints—without requiring a separate engineering team.
3	**Stakeholders Decide**	Clinical, regulatory, and commercial teams explore results and make faster, data-driven decisions on their own terms.

This shift matters because it closes the gap between analysis and action. When a biostatistician can publish a Shiny app directly from the same environment where the model was validated, the time from insight to decision shrinks from weeks to hours.

Platform Architecture: What to Look For

Making this three-stage shift operational requires a platform that can support the full lifecycle—from model development through app publication to ongoing governance. At Domino Data Lab, we've seen this play out across dozens of life sciences organizations, and the architectural requirements are consistent regardless of the specific tools involved.

Unified Development and Deployment

Rather than stitching together separate tools for development, deployment, and distribution, the platform should provide a single environment. Data scientists work in the languages (R, Python, SAS) and frameworks (Shiny, Dash, Streamlit, Flask) they already know, and publish results as interactive apps, APIs, or model endpoints—without a handoff to a separate engineering team. Eliminating that handoff removes a major source of delay and translation loss.

Chain of Custody Built In

Every app should inherit the platform's full traceability: from raw data, through code, to final output. Statistical Computing Environment quality control (SCE QC) workflows, role-based access controls, and version-controlled environments ensure audit readiness at every stage. This is not governance bolted on after the fact—it must be built into the workflow itself.

Reusable Science at Scale

One of the most powerful capabilities is the ability to standardize workflows once and reuse them across studies. A well-designed RWE cohort builder or clinical trial explorer becomes a reusable asset, not a one-off project. Auto-scaling infrastructure ensures apps perform under load when multiple teams access them simultaneously.

Use Cases Across Clinical and RWE

These design and governance principles come alive in three areas where interactive applications are delivering the most tangible impact today.

Clinical Trial Reporting

Interactive tables, listings, and figures (TLFs) allow stakeholders to filter by subgroup, endpoint, and visit without waiting for a programmer to re-run the analysis. Patient-level explorers and safety signal dashboards give clinical teams immediate access to the data they need for decision-making at data monitoring committee meetings and safety reviews.

RWE Insights Dashboards

Real-time cohort selection across claims and electronic medical record data enables teams to explore adherence, treatment patterns, and outcomes with model-integrated analytics. Rather than requesting a new cut of data for each question, medical affairs and HEOR teams can build and refine cohorts interactively, accelerating the evidence generation cycle.

Regulatory Submissions

Governed, reproducible outputs with full chain of custody from data source to final deliverable provide the auditability that health authorities require. When an inspector asks "show me how you got from this dataset to this figure," the answer is embedded in the platform—not reconstructed from memory.

Balancing Innovation with Compliance

A common misconception is that self-service interactivity and regulatory compliance are in tension. In practice, the right platform makes them complementary.

Enabling Innovation

Flexible framework support (Shiny, Dash, Streamlit, and custom options) lets teams choose the best tool for each use case. Self-service exploration with parameterized views allows stakeholders to answer their own questions. Rapid prototyping with pre-built templates accelerates time to value, and model-powered insights can be surfaced directly within the app interface.

Ensuring Compliance

Audit-ready traceability through SCE QC workflows provides a documented, reproducible path from data to output. Role-based access control and data governance ensure that users see only what they are authorized to see. Version-controlled, reproducible environments mean that any result can be recreated exactly as it was originally produced. The full chain of custody—raw data to code to output—satisfies the most demanding regulatory expectations.

The Synthesis

The key insight is that compliance and innovation reinforce each other when embedded in the same platform. Governance that is woven into the development workflow does not slow teams down—it gives them confidence to move faster, knowing that every output is traceable and every environment is reproducible. Enterprise platforms supporting these use cases typically maintain certifications such as ISO 9001, ISO 27001, SOC 2, GxP, HIPAA, and GDPR, providing the compliance foundation that regulated industries require.

Acknowledging the Risks

Interactive applications are not without pitfalls. Self-service filters can lead to misleading subgroup analyses if guardrails are not built in. App proliferation without lifecycle management creates maintenance burden. And if apps are not designed with clear defaults, users may draw conclusions from views the original analysts never intended. Acknowledging these risks early—and building in safeguards like pre-defined filter ranges, usage analytics, and clear methodological documentation within the app—is essential for responsible deployment.

Scaling Apps Sustainably Across Organizations

Building one successful app is a proof of concept. Scaling interactive applications across an organization is a strategic capability. The journey from pilot to platform requires deliberate investment in four areas.

1. Identify high-value use cases. Focus on analyses characterized by repeated questions, multiple audiences, or evolving data. Not every analysis warrants an app, but the ones that do will generate outsized returns in time saved, decisions accelerated, and stakeholder satisfaction.

2. Standardize and template. Create reusable app templates with consistent user experiences and pre-approved components. Standardization reduces development time for each subsequent app and ensures a consistent experience for users across the organization.

3. Build internal capability. Train data scientists and statistical programmers to publish apps directly from their analysis environments. This reduces dependency on engineering teams and empowers the people closest to the data to deliver the final product.

4. Govern and iterate. Establish app lifecycle policies covering versioning, retirement, and ongoing quality control. An app is not a one-time deliverable—it requires maintenance, monitoring, and periodic review just like any other production system.

Measuring Long-Term Impact

Scaling sustainably also means measuring whether apps are delivering value. Track adoption metrics (active users, session frequency), operational metrics (time saved versus static report cycles, reduction in ad-hoc re-analysis requests), and decision metrics (time from data availability to stakeholder action). Reviewing these on a quarterly basis helps teams identify which apps to invest in further, which to simplify, and which to retire.

Conclusion

Interactive applications represent a fundamental shift in how clinical and RWE teams deliver value. When done well, they transform stakeholders from passive consumers of reports into active explorers of data, accelerating the path from insight to decision.

The keys to success are straightforward but require discipline: choose the right use cases, design for the decision rather than the data, embed compliance into the workflow rather than layering it on afterward, and invest in the organizational capabilities needed to scale.

The organizations that get this right will not just move faster—they will make better decisions, produce more trustworthy evidence, and build a durable competitive advantage in an industry where the quality and speed of evidence generation increasingly determines outcomes.

About the Author

Chris McSpiritt is at Domino Data Lab, where he works with life sciences organizations to scale data science and AI capabilities. For questions or to continue the conversation, reach out at christopher.mcspiritt@dominodatalab.com.

About Domino Data Lab

Domino Data Lab provides the leading enterprise AI platform for highly regulated industries. Six of the top ten pharmaceutical companies use Domino to power AI and analytics across discovery, development, and manufacturing. Learn more at domino.ai.