Modernizing the Vaccine TAUG: An End-to-End, Output-Driven Framework for Reactogenicity, Immunogenicity, Efficacy, and Reviewer-Ready Submissions

PharmaSUG China 2026

Lin CHEN, Sanofi, Shanghai, China; Jialiang LIU, Sanofi, Chengdu, China

Abstract

The Clinical Data Interchange Standards Consortium (CDISC) Therapeutic Area User Guide for Vaccines (TAUG-Vax) v1.1 (9 April 2018) remains a useful reference for reactogenicity data, but it was issued as a provisional, Study Data Tabulation Model (SDTM)-oriented guide built on the foundational standards of that period (SDTM v1.4 / SDTM Implementation Guide [SDTMIG] v3.2). Contemporary vaccine programs are far broader: multi-dose and booster schedules, multivalent and co-administration designs, adaptive master protocols, humoral and cell-mediated immunogenicity, disease-endpoint efficacy, and immunobridging. CDISC foundational standards have also moved on — SDTMIG v3.4 revised the Immunogenicity Specimen Assessments (IS), Laboratory Test Results (LB), and Microbiology Specimen (MB) domain scopes and added the Cell Phenotype (CP) and Genomics (GF) Findings domains [2] — and U.S. Food and Drug Administration (FDA) Center for Biologics Evaluation and Research (CBER) Office of Vaccines Research and Review (OVRR) technical expectations have become more specific [3]. As a result, v1.1 is no longer sufficient as an end-to-end implementation guide.

This paper identifies the main gaps in TAUG-Vax v1.1 across reactogenicity, immunogenicity, and clinical endpoints (efficacy and immunobridging), then proposes a modernized, COVID-19-TAUG-style guide organized as an output-driven framework in which implementation decisions are traced from protocol and Statistical Analysis Plan (SAP) through Clinical Data Acquisition Standards Harmonization (CDASH), SDTM, and Analysis Data Model (ADaM) to Clinical Study Report (CSR) tables, registry and lay outputs, reviewer guides, and validation. A worked example follows one solicited injection-site reaction from source data to a reviewer-traceable tables, listings, and figures (TLF) cell. All datasets, variables, and patterns shown are illustrative, generalized, and de-identified; they are based on public CDISC and FDA standards and on SDTM and ADaM principles, not on any sponsor-specific template or controlled document.

Keywords: Vaccine TAUG, reactogenicity, immunogenicity, SDTM, ADaM, traceability.

List of Abbreviations

Abbreviation	Definition
aCRF	Annotated Case Report Form
ADaM	Analysis Data Model
ADADA	Anti-Drug Antibody Analysis Dataset
ADIS	Immunogenicity Analysis Dataset
ADRC	Reactogenicity Analysis Dataset
ADRG	Analysis Data Reviewer's Guide
ADTTE	Time-to-Event Analysis Dataset
AE	Adverse Event
AESI	Adverse Event of Special Interest

Abbreviation	Definition
AP	Associated Persons
ARS/ARD	Analysis Results Standard / Analysis Results Data
BIMO	Bioresearch Monitoring
CBER	Center for Biologics Evaluation and Research
CDASH	Clinical Data Acquisition Standards Harmonization
CDISC	Clinical Data Interchange Standards Consortium
CE	Clinical Event
CHIM	Controlled Human Infection Model
CoP	Correlate of Protection
CORE	CDISC Open Rules Engine
CP	Cell Phenotype Findings domain
CSR	Clinical Study Report
CTG	ClinicalTrials.gov
CTIS	Clinical Trials Information System
DTP	Data Transfer Plan
EAS	Evaluable Analysis Set
eCOA	Electronic Clinical Outcome Assessment
eCTD	Electronic Common Technical Document
eDiary	Electronic diary
EX	Exposure domain
FACE	Findings About Clinical Events domain
FAS	Full Analysis Set
FDA	U.S. Food and Drug Administration
FOCID	Focus of Study-Specific Interest Identifier
GF	Genomics Findings domain
GMC	Geometric Mean Concentration
GMT	Geometric Mean Titer
GMFR	Geometric Mean Fold Rise
GMR	Geometric Mean Ratio
HR	Hazard Ratio
IS	Immunogenicity Specimen Assessments domain
ISS/ISE	Integrated Summary of Safety / Effectiveness
LB	Laboratory Test Results
LLN/ULN	Lower / Upper Limit of Normal
LLOQ/ULOQ	Lower / Upper Limit of Quantitation
MAAE	Medically Attended Adverse Event
MB	Microbiology Specimen

Abbreviation	Definition
NC	Not Calculated
OVRR	Office of Vaccines Research and Review
PIMMC	Potential Immune-Mediated Medical Condition
PPAS	Per-Protocol Analysis Set
RELREC	Related Records dataset
RELSUB	Related Subjects dataset
RR	Relative Risk / Rate Ratio
SAP	Statistical Analysis Plan
SDRG	Study Data Reviewer's Guide
TLF	Tables, Listings, and Figures
VE	Vaccine Efficacy
VLM	Value-Level Metadata
VS	Vital Signs

1. Introduction

Vaccine clinical trials differ from general drug trials in several ways. Beyond standard safety assessments, vaccine studies commonly include solicited reactogenicity events, unsolicited Adverse Events (AEs), Medically Attended Adverse Events (MAAEs), Adverse Events of Special Interest (AESIs), Potential Immune-Mediated Medical Conditions (PIMMCs), immunogenicity endpoints, clinical disease-endpoint efficacy, and complex dose schedules. These data are analyzed by dose, antigen, visit, time window, baseline serostatus, assay method, region, age group, and analysis set — few therapeutic areas combine this many analysis dimensions in one submission.

TAUG-Vax v1.1 [1] provided an important foundation for representing vaccine data, but it was explicitly a provisional guide focused on reactogenicity and vaccine administration, written against SDTM v1.4 and SDTMIG v3.2. It deliberately excluded the CDASH collection model, ADaM construction, unsolicited AEs, PIMMCs, clinical efficacy, and complex immunogenicity. Since 2018, the landscape has shifted substantially with broader use of mRNA and viral-vector platforms, adaptive and co-administration designs, and multi-year durability follow-up. CDISC foundational standards have evolved in parallel [2], and FDA CBER, through OVRR, has issued technical-conformance expectations [3] that refine and, in some instances, go beyond the 2018 examples. v1.1 offers useful reactogenicity examples yet lacks end-to-end direction from SDTM mapping through ADaM derivation to traceable review outputs. This paper consolidates the implementation gaps, corrects several common misconceptions, and proposes a framework for a future guide.

2. Why TAUG-Vax v1.1 Is No Longer Sufficient

Several characteristics limit v1.1's applicability to current studies:

- **Narrow scope.** Centered on reactogenicity safety; modern submissions also need immunogenicity, disease-endpoint efficacy, MAAEs, AESIs, PIMMCs, assay metadata, and cross-regional reporting.
- **SDTM-oriented, not end-to-end.** Insufficient ADaM examples for Geometric Mean (GM), Geometric Mean Ratio (GMR), Geometric Mean Fold Rise (GMFR), seroconversion,

seroprotection, vaccine response, baseline-serostatus subgrouping, or per-protocol immunogenicity analysis.

- **Older foundational standards.** Predates the SDTMIG v3.4 IS/LB/MB scope changes and the new CP domain [2]; unchanged examples lead different sponsors to implement immunogenicity data inconsistently across studies.
- **Limited population/denominator guidance.** Safety and immunogenicity denominators depend on vaccinated status, evaluable diary data, at-risk intervals, baseline serostatus, valid assay results, and visit windows; v1.1 provides no comprehensive ADaM framework for these decisions.
- **Unresolved known issues.** v1.1 acknowledges uncertainty around imprecise assessment intervals, evaluator attribution, event-end dates beyond the observation interval, collected maxima, Clinical Event (CE) summary anchoring, and whether severe or long-lasting reactions belong in AE, CE, or both [1]. These are programming rules that drive SDTM, Reactogenicity Analysis Dataset (ADRC), define.xml, and reviewer documentation.
- **Cross-TAUG inconsistency.** TAUG-Influenza v1.1 [16] remains SDTM-only, while TAUG-COVID-19 v2.0 [15] adds CDASH, Annotated Case Report Forms (aCRFs), SDTM examples, and ADaM guidance. A future Vaccine TAUG should adopt the latter end-to-end structure and harmonize examples across products.

3. Gap 1 — Reactogenicity Mapping Still Needs Stronger Regulatory Alignment

Reactogenicity refers to expected post-vaccination inflammatory reactions, pre-specified in the protocol and collected within a fixed window via paper diary or eDiary/eCOA (often called “solicited AEs”), divided into administration-site and systemic events. It is the strongest area of v1.1, yet implementation has moved well beyond the original examples.

3.1 The Flat Model and Core SDTM Mapping

v1.1 allowed either a nested model (records only when an event occurs) or a flat model. Current OVRR technical guidance expects, or strongly recommends, the flat model [3]: within the assessment window (e.g., 7 or 14 days), every assessed day is represented, and an explicit assessment record is generated even when the subject reports no symptom. This improves reviewability but substantially increases dataset size in large studies, so the updated guide should give decision criteria for flat, nested, or hybrid models and require the choice to be documented in the SDRG. Reactogenicity data map into three core SDTM domains: CE (a global per-subject / per-vaccination / per-symptom occurrence-and-duration summary), FACE (daily severity, characteristics, and results), and VS (quantitative temperature for fever). CECAT is set to 'REACTOGENICITY' and CESCAT to 'ADMINISTRATION SITE' or 'SYSTEMIC'. Summary and daily records should be connected through CE.CELNKGRP with FACE.FALNKGRP / VS.VSLNKGRP via RELREC, while FOCID should preserve anatomical location where co-administration, repeated injection sites, or heterologous boosters could otherwise conflate site-specific reactions.

3.2 Continuing Events: CE/AE Allocation

A central question is how to represent reactions that continue beyond the solicited interval or become severe enough to meet AE-reporting criteria. v1.1 identifies the issue but does not resolve it, and the answer affects whether a reaction appears in safety summaries as a clinical event, an adverse event, or both, duration derivation, MedDRA coding, registry counts. The recommendation depends on jurisdiction:

- **FDA / OVRR implementation.** Under current OVRR technical expectations [3], a solicited reaction that continues beyond the planned assessment interval should also be represented as an

AE, while CE retains the solicited-period global summary and FACE/VS retains the daily detail. CE conveys ongoing status through end-reference variables (CEENRTPT/CEENTPT), CE and AE are linked through RELREC, and duration is expressed separately for the solicited window and the continuation period to avoid double counting.

- **General / global implementation.** Other regions or protocols may differ. The CE/AE rule should be pre-specified from protocol criteria, collection design, severity, duration, and any health-authority agreement, and regional differences should be explained in the SDRG/ADRG. For multi-dose trials, an unresolved event at the next vaccination may have its analysis end date truncated at that vaccination to avoid attributing a continuing signal to the wrong dose.

3.3 Reconciliation of Solicited Events Initially Reported as AEs

In decentralized and large studies, a symptom omitted from the diary may be entered on the AE page when an investigator identifies it at a later visit. If its actual onset falls inside the reactogenicity window and it meets the solicited-event definition, the record should be reconciled rather than left as an unsolicited AE. This is a reconciliation procedure, not a statistical imputation: teams compare the AE start date with the vaccination date and the assessment window, apply category control at the eCRF, and use medical-monitor review to decide whether to reclassify. Because such information is recall- or reconstruction-based rather than a contemporaneous diary entry, OVRR expects it to be identified as reconstructed data (e.g., via a supplemental qualifier such as RECON) [3]. Corresponding FACE/VS records should be created only where day-level detail is genuinely supported, and the deduplication logic should be documented in the SDRG/ADRG.

3.4 Evaluator and Source-of-Assessment Hierarchy

This section sets out the solicited-reaction source hierarchy that other sections refer to. The hierarchy is an illustrative proposed analysis convention rather than a CDISC-mandated rule; it should be pre-specified in the protocol, SAP, or study metadata according to the collection design. v1.1 acknowledges that participant-reported and investigator assessments may complement or conflict but provides no agreed representation. A submission-ready approach preserves all collected source records in FACE/VS with evaluator/source metadata, then selects the analysis-used value in ADRC through a pre-specified hierarchy:

- Prefer investigator assessment when available; use participant assessment only when both investigator fields are absent. If occurrence is 'No' and intensity is missing, derive analysis intensity as 'None' (not missing).
- Consistency rules: if occurrence is missing but a measurement meets Grade 1+, derive occurrence as 'Yes'; if occurrence is 'No' but a measurement meets Grade 1+, query the record, and if uncorrected apply a pre-specified conservative (worst-case) rule; if the form is not completed, treat as missing/not evaluable, never as 'no reaction'.

This hierarchy belongs in metadata, program templates, and reviewer documentation so that SDTM, ADaM, and TLFs do not diverge across studies.

3.5 Timing, Duration, and Collected Summary Results

SDTM-valid data can still be analytically ambiguous in timing. Diary days may reflect calendar days, end-of-day completion, or a protocol-defined interval; a value such as -P1D in --EVLINT does not necessarily mean a strict 24-hour interval anchored to the vaccination time. A future guide should distinguish the intended diary day, the eDiary reporting window, and the biological time since vaccination, and represent them consistently across FACE/VS, ADRC, value-level metadata, and TLF footnotes. Collected summaries (e.g., maximum temperature or redness diameter recorded directly) belong in SDTM with a summary-result qualifier and value-level metadata; derived summaries belong in

ADaM with DTYPE and flags identifying collected, derived, analysis-used, or supportive values.

3.6 Representing Missing Expected Assessments (the “NOT DONE” Pattern)

Once the flat model is adopted, missed diary days must be represented explicitly rather than left as unexplained blanks. This is the programmatic generation of structured, expected 'NOT DONE' records — placeholders that record the absence of an assessment, not imputed clinical results. A standard workflow uses Exposure (EX) to identify vaccinated subjects and vaccination references (--TPTREF), generates a continuous series of expected diary days per reference (durations may differ by event type, e.g., 7 days systemic and 14 days administration-site), merges the expected matrix with actual FACE/VS data, and, where an expected day has no actual record, generates a structured record (e.g., FATESTCD = 'OCCUR', no occurrence result, the relevant --STAT set to 'NOT DONE', with --REASND explaining the status). These records are limited to protocol-predefined reactogenicity events. The matrix must respect discontinuation: cross-checking Disposition (DS), a subject who withdraws on a given day receives 'NOT DONE' records only up to the applicable withdrawal day; later planned days are excluded.

CE occurrence (CEOCCUR) and the null-versus-N rule. CE synthesizes diary, symptom-resolution, and unplanned-assessment data. CEOCCUR = 'Y' if any valid assessment day shows the symptom occurred. CEOCCUR = 'N' only when every expected assessment day is evaluable (non-missing) and all are negative. If at least one expected day is missing and no day is 'Y', CEOCCUR should remain null, with --STAT/--REASND explaining the missingness — a present record of 'N' on other days does not justify inferring 'N' overall. The event start date (CESTDTC) captures the first FACE date (FADTC) marked as occurred and remains blank when the status is entirely 'NOT DONE'. These null-versus-N distinctions are common defects and must be applied consistently in the worked example and in any denominator logic.

Table 1. Evolution of reactogenicity mapping rules across guidance generations.

Implementation aspect	v1.1 / legacy issue	Regulatory or implementation signal	Recommended modernization
Preferred data model	Flexible nested or flat	FDA/CBER prefer the flat model [3]	Decision criteria for flat/nested/hybrid; document in SDRG/ADRG.
Continuing-event allocation	Known issue; CE vs AE unresolved	OVRR expects continuation beyond the window to also appear in AE [3]	FDA/OVRR: CE summary + FACE/VS detail + AE for continuation, linked by RELREC; general: pre-specify by protocol and region.
Event-duration mapping	CESTDTC/CEENDTC represents the global event start and end dates	Duration, days of presence, and continuation can diverge	Separate onset, end, duration, days of presence, and continuation; trace to ADRC and outputs.
AE-to-reactogenicity reconciliation	No standardized procedure	Operational AE capture may include omitted diary symptoms; recall data must be flagged [3]	Reconcile by onset window; flag reconstructed data; document deduplication in SDRG/ADRG.

Implementation aspect	v1.1 / legacy issue	Regulatory or implementation signal	Recommended modernization
Evaluator/source conflicts	--EVAL identifies an evaluator; conflicts unresolved	Ambiguity when participant and investigator assessments differ	Preserve sources; define analysis-used hierarchy in ADRC.
Timing and collected maxima	--EVLINT and NSVs (COLSRT/SUMTYP) limited	Earlier draft discussions flagged --STINT/--ENINT as confusing	Clear timing anchors, value-level metadata, and ADaM flags for collected vs derived summaries.

4. Gap 2 — Immunogenicity Is the Largest Missing Component

The most important limitation of v1.1 is the absence of detailed immunogenicity guidance, even though immune response is often a primary or key secondary endpoint. Common analyses include GM, GMR, GMFR, seroconversion, seroresponse, seroprotection, baseline-serostatus subgroups, response by antigen or serotype, non-inferiority or superiority comparisons, persistence, and booster response.

4.1 From Humoral to Cellular Immunity

Humoral assays such as IgG by ELISA produce titers or concentrations that fit the IS structure cleanly. mRNA and viral-vector platforms have expanded the evidence package to cell-mediated immunity, particularly T-cell responses quantified by ELISpot and multiparameter flow cytometry, where subsets are defined by nested gating on cluster-of-differentiation and functional markers. This high-dimensional, marker-defined output does not fit a structure designed for simple antibody titers, which is precisely why the IS/CP boundary matters.

4.2 Clarifying the IS / CP Boundary (A Common Misconception)

Some teams assume the IS domain is now obsolete and that all immunogenicity data should be moved to CP. That is incorrect and can create conformance problems:

- **IS (Immunogenicity Specimen Assessments)** remains the home for specimen-based immune-response testing — neutralizing or binding antibody titers, anti-drug antibody results, and antigen-induced responses, including certain non-flow-cytometry cellular assays. SDTMIG v3.4 expanded IS with variables such as MSCBCE and the cross-domain BDAGNT, TSTCND, and CNDAGT, reducing overload of --TEST/--TESTCD [2,3].
- **CP (Cell Phenotype Findings),** introduced in SDTMIG v3.4 [2], characterizes cell phenotype, lineage, and function from marker expression in single-cell suspensions and is suited to flow-cytometry phenotyping. Marker-expression-based flow-cytometry phenotyping is generally represented in CP, whereas antigen-induced immune-response measurements, including ELISpot and certain functional cellular assays, are represented in IS. The routing decision should therefore be based primarily on the assessment concept rather than on the assay platform alone. CP complements IS; it does not replace it.

Figure 1 illustrates an explicit IS/CP/LB/MB routing decision tree for vaccine specimen-based data; it is consistent with the CDISC specimen-based-findings decision tree [17], and the updated guide should accompany such a tree with vaccine-specific examples.

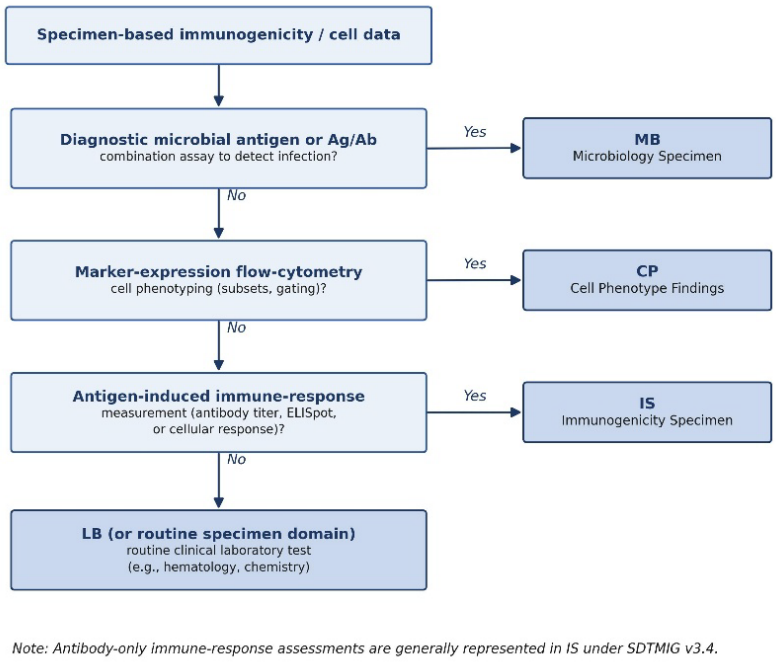

Figure 1. Routing decision tree for vaccine specimen-based immunogenicity and cell data across the IS, CP, MB, and LB domains.

4.3 Mapping Flow-Cytometry Data to the CP Domain

For T-cell phenotyping by flow cytometry, the principal CP variables are: CPCAT (category, typically 'IMMUNOPHENOTYPING'); CPTEST / CPTESTCD (the named cell population); CPSBMRKS and CPCSMRKS (subset- and cell-subset-defining markers that subdivide the population); CPCELSTA (cell state, e.g., 'ACTIVATED'); and CPRESSCL / CPRESTYP (result scale and type, e.g., 'QUANTITATIVE' with 'NUMBER FRACTION' or 'NUMBER CONCENTRATION') [2]. Gating logic is captured through the marker variables and, where additional context is required, supplemental qualifiers; test-mode information is currently carried in the non-standard variable TSTMOD pending harmonization. Because flow cytometry and ELISpot are run by central labs or specialty vendors, transfer formats should be aligned through a Data Transfer Plan (DTP) at project inception. Table 2 shows a representative mapping.

Table 2. Representative external-lab (DTP) to SDTM CP mapping for flow-cytometry data (illustrative).

External-lab concept	Vendor / transfer variable	SDTM CP target	Example / note
Central lab name	LABNAM	(metadata)	Immunology central lab
Subject ID	SUBJID	USUBJID (derived)	Per study convention
Specimen collection date	COLLDT	CPDTC	ISO 8601
Named cell population	POPNAME	CPTEST / CPTESTCD	e.g., effector-memory CD8 subset
Subset-defining markers	SUBPOP	CPSBMRKS / CPCSMRKS	e.g., CDxx+CDyy-

External-lab concept	Vendor / transfer variable	SDTM CP target	Example / note
Cell state	STATE	CPCELSTA	e.g., ACTIVATED
Analytical method	METHOD	CPMETHOD	Flow cytometry
Result and unit	ORRES / ORRESU	CPORRES / CPORRESU	with CPRESSCL / CPRESTYP

ELISpot response results are generally represented in IS rather than CP, because they measure antigen-induced immune responses rather than marker-defined cell phenotypes; only flow-cytometry phenotyping maps to the CP structure shown above.

4.4 Data Flow to ADaM, LLOQ/ULOQ, and ADIS

Future guidance should specify an explicit flow from source assay result through IS/CP to ADaM parameterization, with rules for LLOQ/ULOQ handling, log-transformation, baseline definition, fold-rise derivation, serostatus classification, population flags, and TLF traceability (Table 3). Neutralizing titer must retain antigen, assay method, unit, specimen type, baseline status, and visit relationship rather than being treated as a generic laboratory value.

A commonly used limit-handling pattern (illustrative). The following are illustrative, commonly used analysis conventions rather than universally required rules; the final approach depends on the protocol, SAP, assay validation plan, and any regulatory agreement. In one common pattern, values below LLOQ are analyzed as LLOQ/2, in-range values are retained as observed, and at-or-above-ULOQ values are capped at ULOQ, with duplicate measurements reduced to an individual geometric mean after limit handling. When final thresholds differ from the raw transfer, SDTM preserves the collected result and ADaM implements the documented threshold. Fold-rise derivation is likewise region- and endpoint-aware: a post/baseline ratio using LLOQ/2 may suffice in some settings, while a more conservative approach uses LLOQ (not LLOQ/2) when baseline is below quantitation, to avoid exaggerating fold increases from very small baseline values. Whichever rule is chosen, thresholds and substitution rule should be documented at value level in define.xml and in the ADRG.

The immunogenicity analysis dataset (illustratively ADIS, with anti-drug-antibody analyses aligned to the ADADA structure [8]) is not merely a container for results: it must retain antigen, assay, specimen, visit, baseline definition, treatment reference, and limit-handling context; use separate parameters or value-level metadata for original-scale, log-transformed, fold-rise, and log fold-rise values; and use criterion variables (CRITy/CRITyFL) for seroconversion, seroresponse, seroprotection, vaccine response, and correlate-of-protection thresholds.

Table 3. Common immunogenicity endpoints and their ADaM requirements.

Endpoint	Definition	Primary SDTM source	Key ADaM requirements	Typical denominator
GMT / GMC	Titer/conc by antigen, visit	IS	Original-scale and log-transformed analysis parameters, baseline, AVISIT, and PARAMCD	Evaluable immunogenicity population
GMR	Ratio of group geometric means between treatment groups at a specified visit	IS	Between-group ratio parameter and confidence-interval method	Subjects evaluable at the specified visit in each treatment group

Endpoint	Definition	Primary SDTM source	Key ADaM requirements	Typical denominator
GMFR	Geometric mean of individual post-baseline-to-baseline fold rises	IS	Individual fold-rise parameter, BASE, and summary method	Subjects with both baseline and post-baseline results
Seroconversion	Threshold/fold criterion	IS	CRITy flag, serostatus	Baseline-defined
Seroprotection	Threshold titer	IS	CRITy flag, threshold	Evaluable at visit
Flow-cytometry cell phenotype	Marker-defined cell population, lineage, or state	CP	Cell-population parameter, defining markers, cell state, and relevant gating context	Subjects with a valid phenotyping result
Antigen-induced cellular response	ELISpot or other functional immune-response measurement	IS	Antigen, test condition, secreted molecule or response measure, assay context	Subjects with a valid immune-response assay result

5. Gap 3 — Clinical Endpoints Beyond Reactogenicity: Efficacy and Immunobridging

Beyond solicited reactogenicity and immunogenicity, v1.1 is essentially silent on the clinical-endpoint content that drives many vaccine approvals: disease-endpoint efficacy, broader (unsolicited) safety surveillance, and immunobridging. Clinical-endpoint efficacy studies, with disease as the primary endpoint, have long played an important role in vaccine licensure, including seasonal influenza vaccine development [19], while the COVID-19 development guidance set explicit vaccine-efficacy success criteria [18]. In 2025, FDA leadership described a policy approach for COVID-19 vaccines proposing randomized, controlled clinical-outcome data for lower-risk populations [20]. These examples show why disease-endpoint efficacy and immunobridging should be first-class implementation areas in a modernized guide.

5.1 Vaccine Efficacy as a Case-Driven Endpoint Family

Vaccine efficacy is typically VE = 1 - RR or VE = 1 - HR on confirmed cases between vaccinated and control groups. Unlike continuous immunogenicity endpoints, efficacy is case- and time-driven, depending on a confirmed-case definition, a surveillance period that usually begins after a protocol-defined post-vaccination interval, person-time at risk, and censoring rules. The analysis structure is therefore often time-to-event or person-time based, depending on the estimand, case definition, surveillance window, and recurrence rules. Implementation needs a disease-confirmation pathway linking symptoms, laboratory/virologic confirmation (MB/IS), and adjudication into an analysis-ready structure (illustratively ADTTE) carrying the confirmed-case flag, at-risk start, event/censoring date, person-time, and population. These structures are substantive components of the efficacy analysis rather than optional bookkeeping: FDA CBER guidance has addressed clinical efficacy evidence for COVID-19 and seasonal influenza vaccines [18,19]. For COVID-19 vaccines specifically, the 2025 policy approach described by FDA leadership proposed clinical-outcome efficacy data before approval in lower-risk groups [20]. The guide would benefit from specifying how the case definition, surveillance window, censoring, and recurrence rules map from protocol/SAP into SDTM and ADaM so that VE estimates and confidence intervals are reproducible.

5.2 Unsolicited AEs, MAAEs, AESIs, and PIMMCs

Unsolicited safety is broader than the solicited model and is only named, not structured, in v1.1. The updated guide should describe how unsolicited AEs, MAAEs, AESIs, and PIMMCs are collected, MedDRA-coded, and analyzed, and how they relate to the solicited structure to prevent double counting (see the CE/AE boundary in Section 3.2 and reconciliation in Section 3.3). AESI/PIMMC handling benefits from pre-specified medical concepts and standardized case definitions; the corresponding analysis datasets should carry analysis windows relative to each vaccination, the special-interest classification, time-to-onset, and outcome/seriousness attributes, with denominators consistent with the safety populations used elsewhere. Risk-interval analyses, a recurring expectation for immune-mediated conditions, should be derivable from the same structures.

5.3 Immunobridging and Correlates of Protection

Immunobridging infers vaccine performance by comparing immune responses with those observed in a reference population, formulation, or regimen with established clinical efficacy — for example in variant-adapted boosters where a placebo-controlled efficacy trial is neither feasible nor ethical. An established correlate of protection (CoP) may strengthen this inference but is not always required: GMR-based non-inferiority and seroresponse comparisons can be specified independently of any formal CoP threshold. FDA CBER guidance has recognized immune-response endpoints in defined vaccine-development contexts [19]. For COVID-19 vaccines specifically, the 2025 policy approach described by FDA leadership proposed immunogenicity-based pathways for higher-risk groups, together with clinical-outcome or post-marketing evidence in specified settings [20]. Immunobridging and disease-endpoint structures should therefore be treated as complementary rather than interchangeable. v1.1 does not address this now-central pathway. The data model must link immunogenicity parameters to the bridging comparison, represent the relevant continuous-marker or threshold relationship, and structure datasets so the analysis is reproducible. In practice this means consistent antigen/assay/visit parameterization in ADIS (Section 4.4), a clear bridging-population definition (e.g., between age groups, formulations, or a reference and a variant-adapted product), and — as a related but distinct analysis concept — any CoP threshold documented at value level (CRITy/CRITyFL). A worked immunobridging case — showing how a GMR-based non-inferiority comparison, and separately a CoP threshold where one exists, are derived and traced from IS source through ADIS to the final inference — would be a valuable addition to TAUG-level examples.

Table 4. Disease-endpoint efficacy, unsolicited safety, and immunobridging — SDTM-to-ADaM requirements (illustrative).

Endpoint area	Operational definition	Primary SDTM source	Key ADaM requirements
Vaccine efficacy (VE)	1 - RR or 1 - HR on confirmed cases over a surveillance window	Symptom diaries, MB/IS confirmation, adjudication	Time-to-event or person-time analysis structure: confirmed-case flag, at-risk start, event/censoring date, person-time, recurrence rule, and analysis population
MAAE	AE leading to a medically attended visit	AE with contact/visit qualifiers	Occurrence dataset; medically-attended flag; window relative to vaccination
AESI / PIMMC	Pre-specified special-interest / immune-mediated conditions	AE plus special-interest classification	Special-interest flag, risk-interval window, time-to-onset, outcome/seriousness

Endpoint area	Operational definition	Primary SDTM source	Key ADaM requirements
Immunobridging	Comparison of immune responses with an efficacious reference population, formulation, or regimen	IS (antigen, assay, visit)	Antigen/assay/visit parameterization, GMR or seroresponse comparison, non-inferiority criteria, bridging-population definition
Correlate of protection	Relationship between an immune marker and clinical protection	IS plus the relevant clinical endpoint source	Continuous-marker or threshold parameter, clinical-endpoint linkage, CRITy/CRITyFL where applicable

6. An End-to-End, Output-Driven Modernization Framework

A modern TAUG should be output aware. v1.1 mainly describes where data may be placed in SDTM, but real submissions are driven by a controlled chain from protocol and SAP through SDTM and ADaM to CSR tables, listings, registry and lay outputs, reviewer guides, and validation. When this chain is not specified early, technically valid SDTM can still produce ambiguous denominators, inconsistent analysis-used records, and avoidable reviewer questions. The COVID-19 TAUG v2.0 shows that a therapeutic-area guide can span collection to analysis [15]; the next Vaccine TAUG could adopt the same layered, output-driven structure (Figure 2). The patterns below are generalized and are not company-specific requirements.

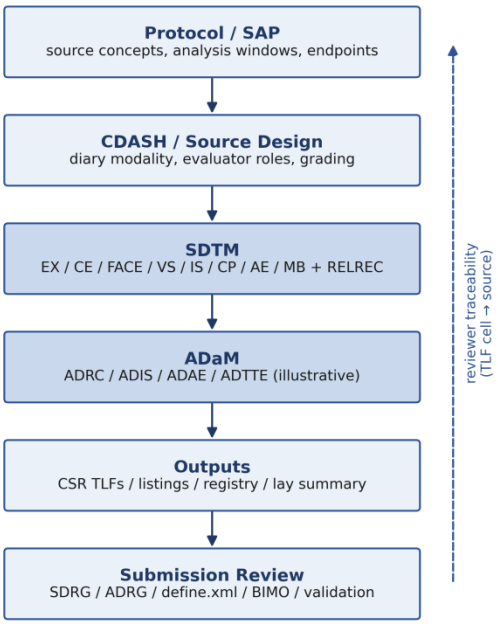

Figure 2. End-to-end, output-driven modernization framework, with reviewer traceability from a TLF cell back to source.

6.1 Source-Concept, Protocol, and CDASH Layer

Define key source concepts (vaccination events, reactogenicity, unsolicited AEs, MAAEs, AESIs, PIMMCs, immunogenicity assays, disease-endpoint cases, concomitant medications, analysis windows) and convert them into collection-ready metadata before SDTM mapping: diary modality, evaluator roles, paper/eDiary reconciliation, exact assessment windows, completion status, global versus daily forms, toxicity-grading scales, continuation follow-up, and device/product attribution, with sample CDASH metadata and aCRF annotations.

6.2 SDTM Layer

Provide updated examples aligned with current SDTMIG scopes, including decision trees for CE, AE, FACE, VS, CM, IS, CP, MB, LB, PE, HO, DD, and RELREC. The reviewer should be able to follow a vaccination in EX to the corresponding FACE/VS daily findings, the CE summary, any AE record, and the analysis-used record without guessing the keys; FOCID/TAETORD are used where they add value (co-administration, repeated sites, product attribution) rather than mandated everywhere.

6.3 ADaM Layer

Naming note: the dataset and variable names used in this paper (ADRC, ADIS, ADTTE, and the variable conventions in the tables) are illustrative analysis conventions consistent with ADaM principles; they are not CDISC-required standard names. A future TAUG should provide ADaM examples for vaccine-specific ADSL flags; reactogenicity derivations integrating FACE daily findings and VS temperature into analysis grades (labeled with DTYPE and traceable to CE/FACE/VS link groups); immunogenicity parameterization (Section 4.4); disease-confirmation/efficacy time-to-event structures; analysis windows relative to vaccination; antigen-level parameterization; and analysis-results metadata for key TLFs, with traceability documented in define.xml and the reviewer guides. ADaMIG v1.3 and ADaM v2.1 [7], the COVID-19 TAUG analysis examples [15], and the Pharmaverse admiralvaccine package [13] provide practical starting points.

6.4 Output and Denominator Layer

The denominator is rarely a single study-wide N. Safety outputs may use a dose-specific safety set, the overall vaccinated population, or participants with available endpoint data; immunogenicity outputs may use the Full Analysis Set (FAS), Evaluable Analysis Set (EAS), or Per-Protocol Analysis Set (PPAS), baseline-serostatus subsets, or subjects with valid assay results at both baseline and post-baseline. Display standards distinguish N(number of participants in the defined population set), M(number of participants with available data for the relevant endpoint), and n(number of participants experiencing the endpoint), and each immunogenicity table should present one analysis set to avoid ambiguity in N; the metadata must define when 'NC' is shown (e.g., single-group geometric-mean CIs suppressed when M is too small or the standard deviation is zero). A dose-level treatment spine (planned/actual treatment, dose number, last dose, co-administered product, route, laterality, injection-site context) should be carried consistently from ADSL into the analysis datasets, listings, and disclosure outputs. Reactogenicity is an endpoint family, not a single flag occurrence (Table 5).

Table 5. Output-driven reactogenicity endpoint family traceable from source to TLF (representations illustrative).

Endpoint	Operational definition	Suggested ADaM / metadata representation
Daily intensity	Per day: None, Grade 1–3 or Missing; measurable reactions use protocol grading thresholds.	Per-day analysis-grade variables derived from FACE/VS, traceable to source link groups.

Endpoint	Operational definition	Suggested ADaM / metadata representation
Maximum intensity	Worst observed grade among evaluable days in the defined period (Grade 3 > Grade 2 > Grade 1 > None); missingness is handled separately through the pre-specified endpoint-evaluability rule and is not ranked below 'None'.	Worst-grade summary variables with an analysis-used flag identifying the record.
Presence	Derived from maximum intensity only when the assessment window meets the pre-specified evaluability rule: Grade 1+ = present; None = absent; otherwise missing/not evaluable.	Presence flag with a denominator based on the pre-specified evaluability rule.
Time to onset	First day with Grade 1+; missing if prior missing intensity prevents reliable determination.	Analysis onset day/category with visit/dose window metadata.
Duration	Last day with presence minus first day with presence plus one; categorized into ranges.	Duration value, unit, and duration category.
Days of presence	Count of days with Grade 1+ during the defined period.	Days-of-presence count and category.
Ongoing status	Assessed around the first day after the solicited period, with rules for missing values.	Ongoing-status flag with end-reference explanation (CEENRTPT/CEENTPT).
Diary completeness	Whether entries were completed within the intended reporting window.	Diary-completion / compliance variables and outputs.

6.5 Submission, Traceability, and Disclosure Layer

Reviewability depends on a self-consistent eCTD package. Many avoidable issues are metadata-level, not clinical-data defects: a broken define.xml hyperlink, an unsynchronized study tagging file, or malformed bookmarks. Teams should reconcile Pinnacle 21 or CDISC CORE findings, submission-portal results, and PDF validation before assembly [4]. Where a BIMO package applies, site-level data should be derived from governed metadata and described in a BIMO Data Reviewer's Guide, consistent with the same population flags and TLF metadata used in the main submission [6]. Listings function as the reviewer's audit trail (global and daily solicited listings, immunogenicity listings, and separate not-in-analysis listings), so listing-ready variables and sort keys should be defined, not only SDTM domains. Finally, the same analysis data supports ClinicalTrials.gov, EudraCT/CTIS, and lay summaries; registry and lay-summary outputs require simplified but internally consistent safety displays aligned with CSR tables and ADaM metadata from the start [5].

6.6 Machine-Readable Content and Validation Layer

A guide delivered only as prose forces each team to reinterpret it. We suggest the guide also be machine-readable, aligned with CDISC Biomedical Concepts and Dataset Specializations and companion controlled-terminology subsets [11], with the output layer aligned to the Analysis Results Standard and Analysis Results Data so the same metadata can drive both static TLFs and interactive review [9]. This direction is consistent with the broader CDISC end-to-end automation program (360i) and the open-source CDISC Open Rules Engine (CORE), which make a 'guide as executable rules' model practical rather than aspirational [10,12]. Value-level metadata(VLM) should manage antigen-by-assay complexity (a small set of parameter-code roots qualified by antigen, method, and unit, with VLM

describing valid combinations, LLOQ/ULOQ, and transformation rules) rather than proliferating --TESTCD/PARAMCD values. An embedded validation catalogue should check the flat-model matrix, 'NOT DONE' correctness, the null-versus-N rule, RELREC linkage, reconciliation logic, multi-dose truncation, LLOQ/ULOQ and serostatus derivations, antigen-by-assay parameter consistency, denominator reproducibility, and TLF-to-ADaM tracing. 7. Complex-Design and Conditional Extensions

Modern vaccine studies add design features that v1.1 barely addresses. Table 6 summarizes these as conditional content; a future TAUG should mark which elements are core and which are design-dependent. For maternal immunization, mother–fetus–infant linkage should reuse the CDISC Associated Persons (AP) model and Related Subjects (RELSUB) dataset [14]. For lot-to-lot consistency, an equivalence comparison of pairwise GMT ratios requires a lot identifier carried from EX into the immunogenicity spine and a documented equivalence-bounds and CI convention.

Table 6. Conditional design extensions and minimum TAUG content.

Extension	Why it matters	Minimum TAUG content
Vaccine co-administration	Discrete products, sequence, site, and a non-inferiority immune-response comparison vs separate administration.	Trial-design (TE/TA) and EX examples preserving product/sequence; treatment-group and attribution rules in SAP and ADaM.
Adaptive platforms and ongoing studies	Rolling submissions and durability follow-up need data-cut-aware logic.	Use ONGOSIND in TS; cut-off-aware DM/DS so status and disposition reflect the cut, not missingness.
Maternal immunization	Mother–fetus–infant linkage and prenatal exposure timing.	Reuse Associated Persons (AP) / RELSUB; worked linkage example [14].
Injection-site reactions	Endpoint-specific onset/persistence/size-grading rules.	Injection-Site Reaction processing separate from systemic AE logic for duration, severity, and grade translation.
Lot-to-lot consistency	Equivalence on pairwise GMT ratios for licensure.	Lot identifier from EX; lot grouping; equivalence bounds and CI method in SAP/ADRG.
Combination products / devices	Pre-filled syringes and devices create product/site/lot/device attribution questions.	A device/product-attribution crosswalk linking EX, device info, product complaints, CE/FACE/VS, AE, and analysis datasets (not a replacement for the Device IG).
CHIM studies	Challenge exposure and protocol-defined infection endpoints.	Capture challenge timing/dose/agent, surveillance window, and endpoint, distinguishing challenge-related from background AEs.
Cross-study poolability (ISS/ISE)	Integrated summaries are easier when designed in from study one.	Consistent parameter codes, denominators, analysis windows, and serostatus rules across studies.

7. Practical Recommendations for Statistical Programmers

- Treat v1.1 primarily as a reactogenicity and vaccine-administration reference, supplemented by current foundational standards, OVRR guidance, and documented implementation considerations.
- Pre-specify the CE/AE boundary for solicited reactions and document it in the SDRG/ADRG (FDA/OVRR: continuation also in AE, linked to CE).
- Document the chosen reactogenicity model (flat/nested/hybrid) in the CDASH design, SDTM specification, and reviewer guide.
- Preserve source records separately from analysis-used records, traceable to the analysis-used value through the Section 3.4 hierarchy.
- Define immunogenicity parameter structures before programming: antigen, assay, visit, baseline, limit handling (Section 4.4), and endpoint derivation across IS/CP, ADaM, and TLFs.
- Treat efficacy and immunobridging as first-class areas: confirmed-case definition, surveillance window, censoring, and CoP thresholds (Section 5).
- Build the 'NOT DONE' generation and discontinuation cross-checks as reusable, validated components, with the null-versus-N rule for CEOCCUR.
- Give each key table a documented denominator rule linked to ADaM flags and source data, and treat TLFs, listings, registry outputs, and lay summaries as design inputs with traceable lineage.

8. Limitations and Scope

This paper is a standards-modernization proposal, not CDISC-endorsed guidance. Accordingly, the framework, conventions, and recommendations here are offered as suggestions for community and CDISC consideration rather than as prescriptions, and any 'should' is intended in that spirit. The worked example is constructed and de-identified, and the dataset and variable names (ADRC, ADIS, ADTTE, and the variable conventions in the tables) are illustrative analysis conventions consistent with ADaM principles, not CDISC-required standard names; the logic does not depend on any particular naming. Regional content is not an exhaustive regulatory comparison, and final LLOQ/ULOQ handling, seroresponse and seroconversion definitions, and immunobridging rules ultimately depend on the protocol, SAP, assay validation, and health-authority agreement. The focus is programming and submission implementation, not clinical or statistical methodology. SDTM/ADaM variable usage should always be confirmed against the applicable published versions for a given submission.

9. Discussion

Modernizing the Vaccine TAUG has direct implications for programming efficiency, reviewability, and submission quality. The most difficult vaccine-standard problems are not limited to domain choice; they involve unresolved business rules at the junction of protocol design, data collection, SDTM representation, ADaM derivation, output denominators, and reviewer documentation. When these rules are not standardized, each sponsor creates local conventions, and reviewers must reconstruct the intent from the SDRG, ADRG, define.xml, and TLF footnotes. Immunogenicity cannot be standardized at the test-name level alone: analytic meaning depends on antigen, assay, unit, limit handling, baseline status, visit window, and endpoint definition. Efficacy and immunobridging add a further layer, tying immune markers to clinical inference through case-driven, traceable structures that v1.1 does not provide. A guide that connects clinical data management, biostatistics, statistical programming, medical writing, pharmacovigilance, regulatory operations, trial transparency, and reviewers — and that is also machine-readable — would be more useful today and more resilient to future platform and design innovation.

The challenges described in this paper are not unique to any single organization. In March 2023, the Vaccines Industry Standards Group (VISG) was established as a cross-industry initiative to address

precisely these kinds of challenges. VISG brings together representatives from leading vaccine manufacturers — AstraZeneca, GSK, Johnson & Johnson, Merck, Moderna, Pfizer, Sanofi — to facilitate open dialogue on the practical difficulties encountered by different sponsors in implementing CDISC standards and meeting regulatory requirements. The alignment between the challenges discussed within VISG and those presented in this paper underscores the industry-wide nature of these issues and highlights the urgent need for a more collaborative and standardized approach. Incorporating the perspectives and collective experience of VISG members into the future development of the Vaccine TAUG could significantly accelerate progress toward more consistent, efficient, and submission-ready data standards across the vaccine industry.

10. Conclusion

v1.1 provided an important foundation for vaccine reactogenicity data, but its known issues map directly to current risks: CE/AE boundary decisions, evaluator conflicts, imprecise diary timing, summary-result representation, cross-TAUG inconsistency, and missing CDASH, ADaM, efficacy, and immunobridging guidance. The next Vaccine TAUG should answer more than 'Which SDTM domain?' It should also specify how the protocol collects the data, how CDASH describes it, how SDTM preserves the source, how the analysis-used value is selected, how the output denominator is constructed, and how the reviewer verifies the lineage. For statistical programmers, that shift means clearer implementation pathways, less rework, greater cross-study consistency, and a package robust for both traditional and metadata-driven review.

Conflict of Interest

This work was funded by Sanofi.

The author is a Sanofi employee and may hold shares or stock options in the company.

References

1. CDISC. Therapeutic Area Data Standards User Guide for Vaccines (TAUG-Vax), Version 1.1 (Provisional). Clinical Data Interchange Standards Consortium, 9 April 2018.

2. CDISC. Study Data Tabulation Model Implementation Guide: Human Clinical Trials (SDTMIG), Version 3.4. Clinical Data Interchange Standards Consortium, 29 November 2021.

3. U.S. Food and Drug Administration, Center for Biologics Evaluation and Research (CBER), Office of Vaccines Research and Review (OVRR). Submitting Study Datasets for Vaccines to the Office of Vaccines Research and Review: Guidance for Industry; Technical Specifications Document. Version 2.1, December 2019.

4. U.S. Food and Drug Administration. Study Data Technical Conformance Guide: Technical Specifications Document. June 2026.

5. U.S. Food and Drug Administration, Center for Drug Evaluation and Research, Office of New Drugs. Standard Safety Tables and Figures: Integrated Guide. Version 2.0, April 2025.

6. U.S. Food and Drug Administration, Center for Drug Evaluation and Research. Bioresearch Monitoring Technical Conformance Guide: Technical Specifications Document. Version 3.1, September 2024.

7. CDISC. Analysis Data Model Implementation Guide (ADaMIG), Version 1.3, and Analysis Data Model (ADaM), Version 2.1. Clinical Data Interchange Standards Consortium, 2021.

8. CDISC. ADaM Implementation Guide for Anti-Drug Antibodies (ADADA), Version 1.0. Clinical Data Interchange Standards Consortium.

9. CDISC. Analysis Results Standard (ARS) and Analysis Results Data (ARD). Clinical Data Interchange Standards Consortium.

10. CDISC. CDISC 360i: End-to-End Standards Automation Program. Clinical Data Interchange Standards

Consortium.

11. CDISC. Biomedical Concepts and Dataset Specializations (CDISC Library). Clinical Data Interchange Standards Consortium.

12. CDISC. CDISC Open Rules Engine (CORE): Open-Source Conformance Rules Engine. Clinical Data Interchange Standards Consortium.

13. Pharmaverse. admiralvaccine: Vaccine Extension Package for ADaM in R. Package documentation and vaccine-specific ADaM templates.

14. CDISC. Study Data Tabulation Model: Associated Persons (AP) Implementation Guide, and Related Subjects (RELSUB) materials. Clinical Data Interchange Standards Consortium.

15. CDISC. Therapeutic Area User Guide for COVID-19, Version 2.0 (Provisional). Clinical Data Interchange Standards Consortium, 2023.

16. CDISC. Therapeutic Area User Guide for Influenza, Version 1.1. Clinical Data Interchange Standards Consortium.

17. CDISC. Where Does My Lab Data Go in SDTMIG v3.4? (Specimen-based Findings decision tree). CDISC Knowledge Base article.

18. U.S. Food and Drug Administration, Center for Biologics Evaluation and Research. Development and Licensure of Vaccines to Prevent COVID-19: Guidance for Industry. June 2020.

19. U.S. Food and Drug Administration, Center for Biologics Evaluation and Research. Clinical Data Needed to Support the Licensure of Seasonal Inactivated Influenza Vaccines: Guidance for Industry. May 2007.

20. Prasad V, Makary MA. An Evidence-Based Approach to Covid-19 Vaccination. New England Journal of Medicine. 2025;392(24):2484-2486. DOI: 10.1056/NEJMsb2506929.