



Sycamore Clinical Agentic Framework™

SPD 
PROTOCOL

MDR 
METADATA

DTM 
TRANSFORM

DSW 
OPEN-SOURCE

SCE 
ANALYSIS

CDR 
REPOSITORY

AI-Assisted SDTM Transformation

Where Automation Works — and Where Humans Still Matter

6th August 2026 - HARI GOPICHANDRAN & DAKSH SINGH

Practical governance for scaling SDTM mapping with generative AI —
without compromising compliance or reviewer confidence.



About the Presenters



Harinarayan, Gopichandran is an executive with 23+ years of progressive leadership at the intersection of Biostatistics, Statistical Programming, and Clinical AI - spanning pharmaceutical giants (Novartis, Allergan, AstraZeneca, BMS) and cutting-edge life sciences technology firms (Sycamore Informatics, Saama, Cognizant).

His career has been defined by one core mission: accelerating drug development through smarter, technology-driven clinical data solutions - from hands-on statistical programming to architecting enterprise-grade AI frameworks.

Currently, as Sr. Director, Product Services at Sycamore Informatics, he leads deployment of the Clinical Agentic Framework™ - a next-generation AI platform transforming how pharmaceutical companies execute clinical data operations.

He believes the future of drug development belongs to those who can bridge deep domain expertise with emerging AI and he is actively building that future.



Sycamore Clinical Agentic Framework™

SPD MDR DTM DSW SCE CDR
PROTOCOL METADATA TRANSFORM OPEN-SOURCE ANALYSIS REPOSITORY



Hari.Gopichandran@



Daksh Singh is a Business Analyst at Sycamore for almost 2 years. He is a dynamic problem-solver and technology professional based in Indianapolis. Holding a bachelor's degree in computer science from Purdue University, Daksh brings a strong foundation in computational thinking and a passion for rapid, effective execution to the technology and informatics sectors.

In this role, he is deeply involved in executing seamless client implementations and spearheading innovative projects centered around agentic AI. By effectively bridging the gap between technical potential and practical business needs, Daksh ensures the successful deployment of cutting-edge, intelligent solutions for clients.



Sycamore Clinical Agentic Framework™

SPD MDR DTM DSW SCE CDR
PROTOCOL METADATA TRANSFORM OPEN-SOURCE ANALYSIS REPOSITORY



Daksh.Singh@



Current SDTM Transformation Is Hitting Its Limits

01

Time Consumption

The end-to-end process is notoriously slow, creating bottlenecks across the clinical trial reporting pipeline.

02

Intensive Manual Effort

Highly repetitive, manual mapping work is prone to human error and reviewer fatigue.

03

Complex Study Designs


Adaptive, non-standard trial designs make mapping novel data structures increasingly difficult.

04

Evolving Standards

Keeping pace with CDISC updates, regulatory expectations, and sponsor-specific rules demands constant rework.

AI-Powered SDTM Skill — A Human-Checked Pipeline



H = Human checkpoint — accept, modify, or reject before the pipeline advances

Where Humans Still Matter

AI HANDLES

- Standard, previously-mapped domains
- High-volume, repetitive transformations
- Pattern matches against CDISC IG rules
- Draft SAS code from approved mappings
- First-pass P21 / IG compliance checks

HUMAN HANDS OFF TO

- Ambiguous SUPPQUAL mappings
- Sponsor-specific custom domains
- Adjudicated safety data
- First-in-class trial designs with no library precedent
- Final accept / modify / reject decision at every stage

GxP-Compliant Clinical Agentic Framework

Governed by the 4 C's — the system cannot operate as a black box.



Context

Reads and understands the organization's specific standards, custom macros, and historical data structures.



Constraint

No external communication; strict data access controls. Critical actions use deterministic, non-LLM pathways where generative AI isn't needed.



Control

“Human in the driver's seat.” Full transparency into LLM activity at all times; only whitelisted, safe models are used.



Compliance

Total auditability — full activity trail, strict version control, and end-to-end traceability from raw data to final dataset.

Every Suggestion Comes With a Rationale

When the skill proposes a mapping, it surfaces why — so reviewers aren't asked to trust a result they can't trace.

PROPOSED MAPPING — SUBJECT REVIEW QUEUE

Source Field

Proposed Target

Rationale — Precedent

Rationale — Standard

Accept · Modify · Reject

Lab_Value (raw lab extract)

LB.LBORRES → LB.LBSTRESN (after unit standardization)

Matches accepted mapping pattern from Study X (98% field-name similarity)

Consistent with CDISC SDTM IG rule LB.7 for numeric result variables

Unit token matched against controlled terminology (UNIT codelist)

Continuous Learning: The Feedback Loop

PROPOSE

For missing or novel data points, the AI generates proposed mappings and puts them forward for expert review.

REVIEW

The human user maintains total control to accept, modify, or reject these suggestions.

LEARN

Once a mapping is accepted, it is instantly fed back into the organization's standard library.

SCALE

Because of this loop, future transformations become progressively smarter, faster, and uniquely aligned with sponsor-specific nuances.

Time Efficiency Without Sacrificing Quality

8–12 Weeks → 70% Reduction

Historical timeline to get complex-study SDTM datasets ready, compared to the AI-assisted skill.

~60.0% TIME SAVED

SDTM results from a live, multi-week sponsor preview engagement — matched traditional vs. skill-assisted task hours (see next slide).

Resource Reallocation

Programmers and data managers shift from repetitive manual mapping to high-value review, QA, and complex anomaly handling.

Scalability

New standards and previously-mapped datasets are digested automatically, letting trial volume scale without linear headcount growth.

A Progressive Delivery Model, Proven in a Live Preview

Each independently tested and signed off by the sponsor's own team before the next began.



Each week follows the same 5-day model: Days 1–2 (demo + shadowing), the sponsor team drives Days 3–5 (independent testing + structured close-out and sign-off).

SKILL	STATUS	RESULT	DETAIL
SDTM Gen	COMPLETE	~ 60%	time saved, matched tasks
SDTM Review	COMPLETE		

Anonymized results from a live, multi-week sponsor preview engagement. Figures reflect matched traditional vs. skill-assisted task hours only.

ROI SUMMARY

Quantifying the SDTM Gen & Review

Skill tested and metrics validated. Figures reflect only tasks with both Traditional and Skill hours logged (apples-to-apples measurement).

[~60%]

time saved (matched tasks only) on
SDTM Gen and SDTM
Review/Qual

+72.5 hrs of traditional effort excluded – measured against the skill

1

Traditional effort: 131.7 hrs

2

S-CAF Skill effort: 53.5 hrs

3

Total saved: 78.2 hrs

4

TFL Gen: Awaiting review (ROI TBD)

End-to-End Traceability

RAW DATA ORIGINS

End-to-end traceability detailing exactly where source data originated before mapping occurred.

VERSION CONTROL GOVERNANCE

Enforces strict version control across organizational standards, SDTM domain programs, and output datasets to prevent unapproved changes.

SAS CONFIGURATIONS & DEPENDENCIES

Captures the specific SAS configurations utilized and permanently documents all dependencies, including specific custom macros.

COMPREHENSIVE AUDIT TRAIL

Maintains an immutable log capturing all user actions and LLM activities to ensure auditability for regulatory review.

Where Do Regulators Stand on AI-Assisted Submissions?

No explicit prohibition — sponsor accountability is unchanged.

Current FDA and PMDA posture does not bar AI-assisted data transformation in submissions. The sponsor remains fully accountable for data quality, traceability, and validation — which is exactly what the 4C framework and human-in-the-loop checkpoints are designed to preserve.

Full audit trail from raw data to final dataset

Human review at every pipeline checkpoint

Deterministic pathways for critical, non-generative steps

Version control and Traceability across standards, programs, and datasets

BUILT ON THE 4 C'S

Context. Constraint. Control. Compliance.

AI accelerates the repetitive work — humans stay accountable for judgment, ambiguity, and final compliance sign-off.

Questions & Discussion

[Presenter Name] · [Email] · Sycamore Informatics