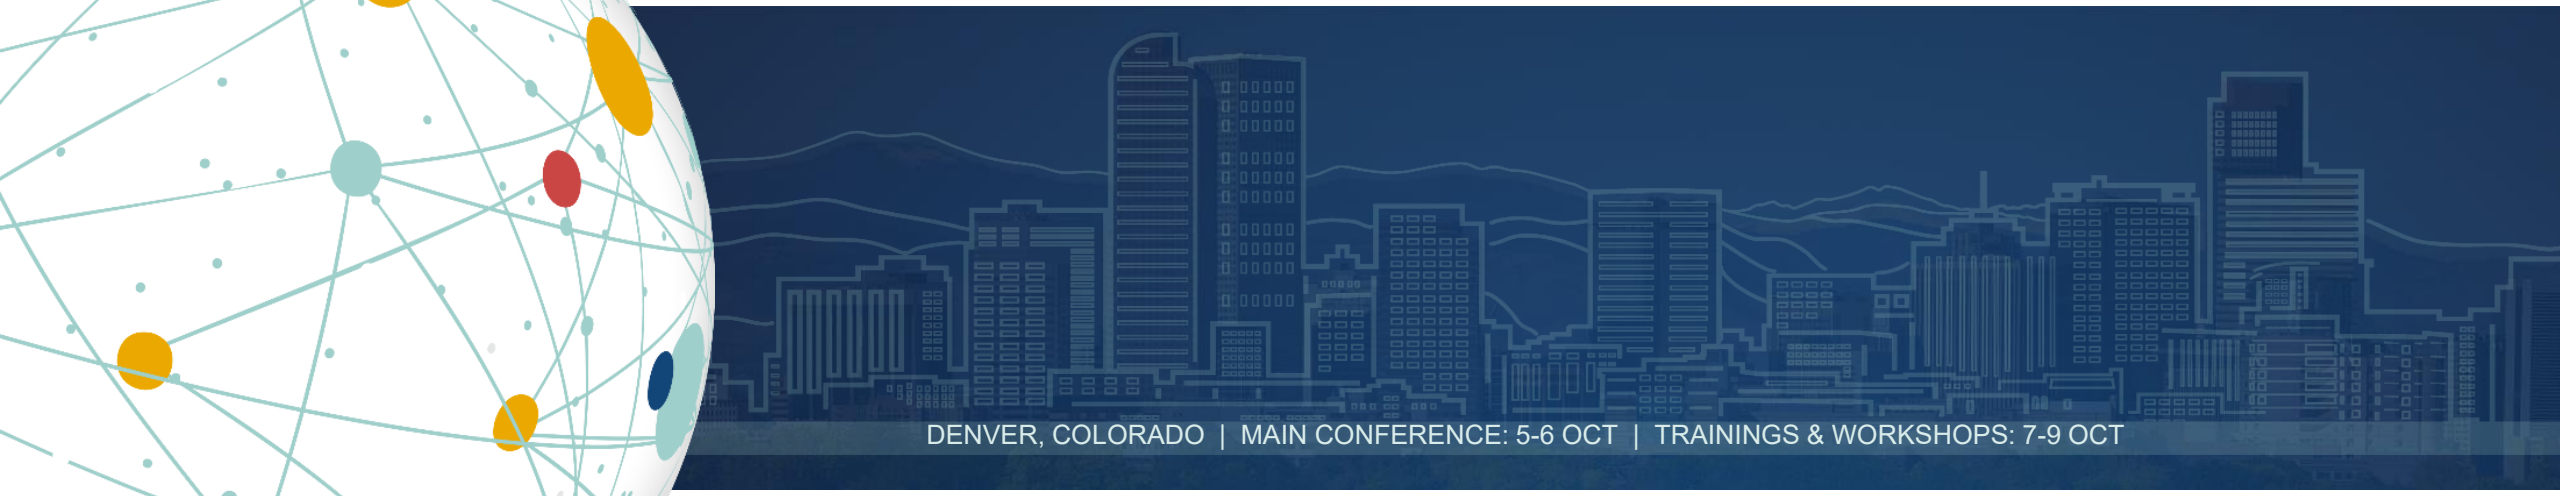


cdisc®

2026 US Interchange

THE FUTURE IS CONNECTED: STANDARDS AND AI POWERING DIGITAL TRANSFORMATION



Approaches to AI Validation in Regulated Clinical Systems and the eTMF

A best-practice, inspection-defensible approach to AI in the regulated eTMF

Presented by Jay Smith, Vice President Product, TransPerfect Life Sciences

Speaker



Jay Smith

Vice President, Product

TransPerfect Life Sciences

Jay currently leads product and tech at Transperfect, responsible for the Trial Interactive platform. Prior to Transperfect, Jay has led product teams at Medidata, Sparta Systems, and Thomson Reuters. Jay has supervised the creation and management of eTMF, CTMS, EDMS, QMS, RTSM, EDC, LMS, and Site Portal solutions. Further back, Jay has also created and managed products for RIMS and eCTD submission publishing and review.

Agenda

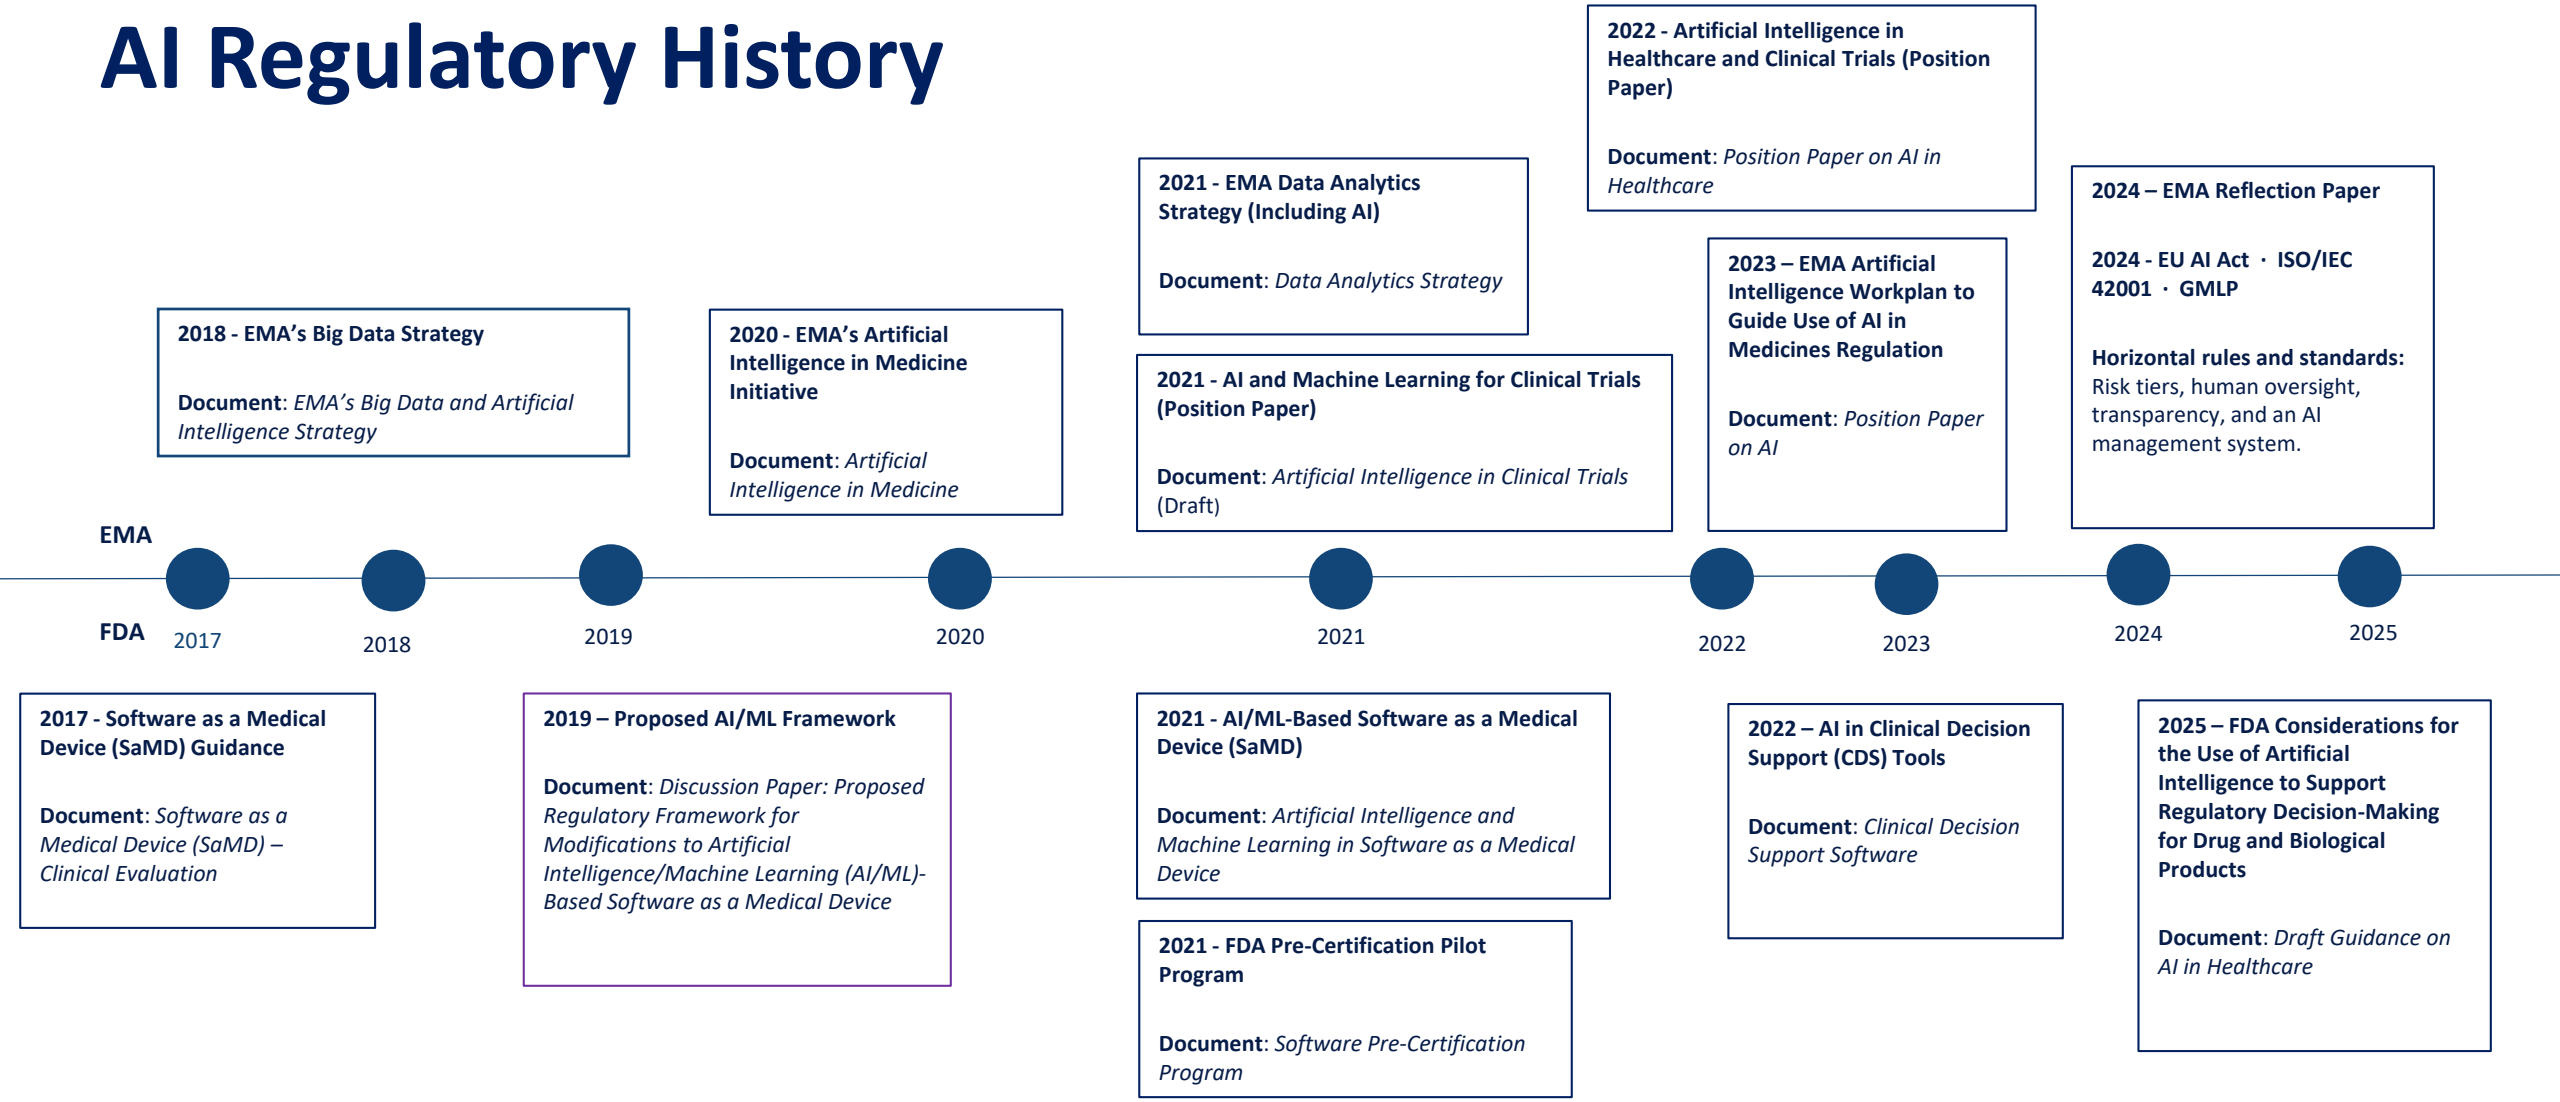
1. Probabilistic vs Deterministic
2. Regulations Timeline
3. Ten Principles
4. Human in the Loop
5. Value vs. Observability
6. How to Validate?
7. The Locked State Approach
8. How to Stay Validated?
9. A Verified Rule-Based Approach
10. TMF Completeness Engine
11. Actions and Next Steps

"Before ML, the software we all worked with was deterministic and always produced the same outputs given the same inputs, following predefined rules. ***Probabilistic*** *software incorporates randomness and statistical likelihood to generate a range of variable but likely outputs for the same inputs.* ***Deterministic*** *software is easy to verify by QA or business unit. AI, and Generative AI in particular, is not."*

- Pratik Daga, Palo Alto Networks, Asana



AI Regulatory History



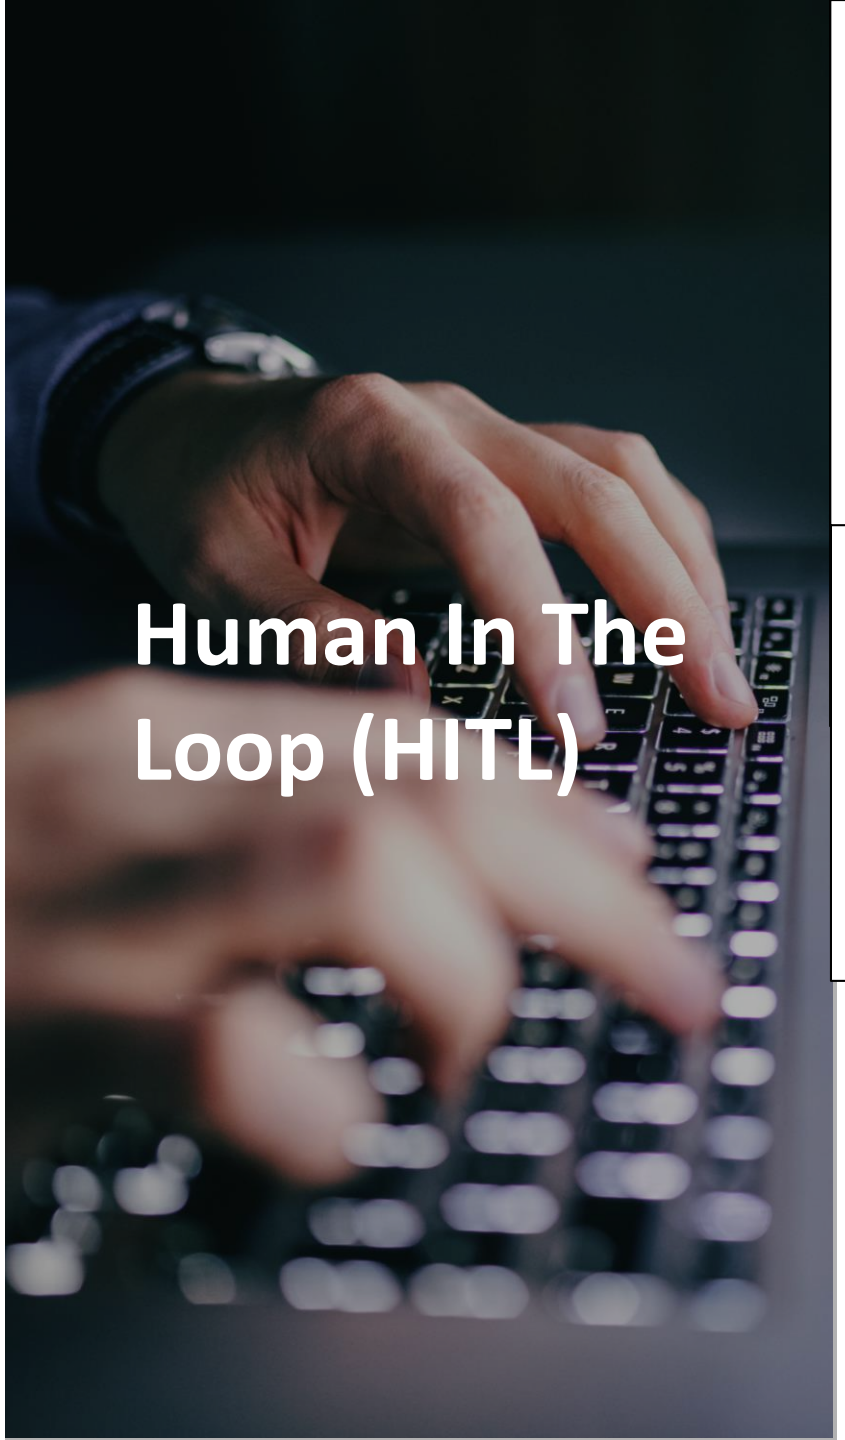
FDA and EMA: Ten Principles of good AI practice

Issued jointly in January 2026, spanning the full drug-development lifecycle.



- Human-centric by Design
- Adherence to Standards
- Multidisciplinary Expertise
- Model Design and Development Practices
- Life Cycle Management
- Risk-based Approach
- Clear Context of Use
- Data Governance and Documentation
- Risk-based Performance Assessment
- Clear, Essential Information

Source: FDA CDER/CBER and EMA, Guiding Principles of Good AI Practice in Drug Development, 14 January 2026.



Human In The Loop (HITL)

1.

Must be explainable and transparent.

2.

Cannot replace human accountability for regulatory compliance.

3.

Should have appropriate risk management strategies.

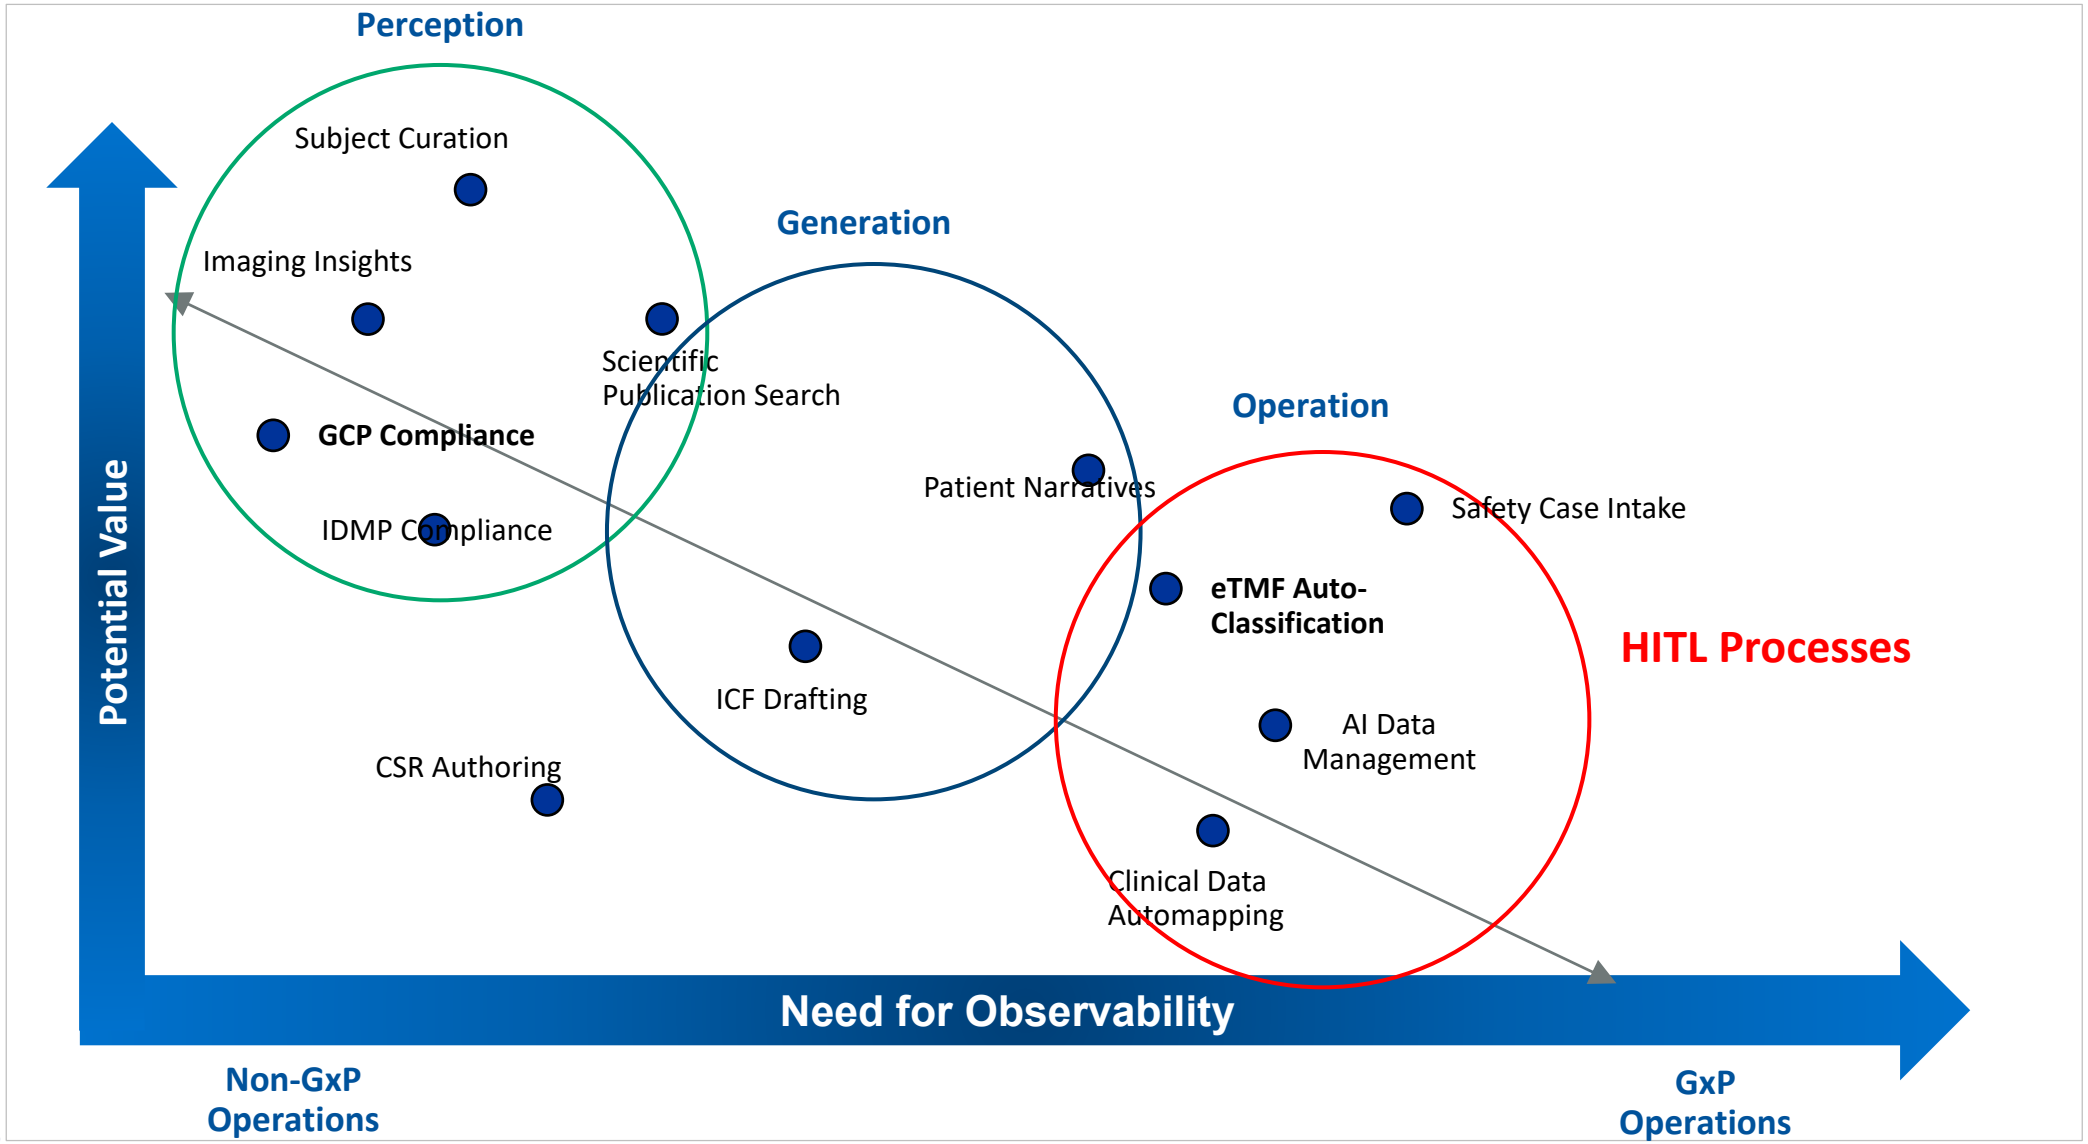
4.

Must comply with GDPR and data protection requirements.

1 https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological

2 https://www.ema.europa.eu/en/news/artificial-intelligence-workplan-guide-use-ai-medicines-regulation

Value vs Observability



The Locked-State Approach for TMF Classification

Allows a probabilistic model to operate as a validated, inspectable system.

Definition. An approach in which the AI model version, the configuration and prompt parameters, the document mappings, the confidence-scoring logic, and the runtime environment are all fixed and version-controlled. No self-learning or autonomous retraining occurs in production.

- Locked-State means controlled so that any change to the model, prompts, thresholds, or mappings goes through formal change control before release.
- Because the model does not learn in production, uncontrolled model drift is structurally impossible.
- Each transaction is assigned a unique fingerprint, recorded in the audit trail for every AI-processed document. That gives reproducibility and full traceability of results.

cdisc®

Source: Proposed Validation Plan for TransPerfect Automate, Section 14.4.

How does a locked state satisfy regulators?

No Uncontrolled Drift	No autonomous retraining in production; every change runs through change control.
Reproducibility and Traceability	A configuration fingerprint on every processed document ties each output to a validated version.
Risk-based Performance Assessment	Each release is validated against predefined acceptance criteria before it is deployed.
Clear Context of Use	The intended-use scope is fixed and validated, not open-ended.
Human-Centric by Design	Confidence thresholds route uncertain outputs to a person; humans stay accountable.

cdisc®

Validating a locked-state AI: the lifecycle

A GAMP 5 risk-based lifecycle, aligned with Annex 11 and the FDA credibility framework.

1. **Intended Use & Context of Use**
2. **Risk Assessment**
3. **Dataset Governance**
4. **Performance Qualification**
5. **Confidence & Human Oversight**
6. **Audit Trail & Dashboard**
7. **Change Control**
8. **Periodic Review & Monitoring**

Validation is continuous. The loop from monitoring back into change control is what keeps the system in a validated state.

cdisc®

Example ‘Locked State’ Verification

Library Analytics

Library List / Library Dashboard: 01. Clinical Trials ET...

Download Report

Latest Run Status

Run ID: f5bb78...b6e3f5

Date: 05/05/2026

Macro F1 Score

96.8%

Threshold: ≥85%

Lowest Class: 94.7%

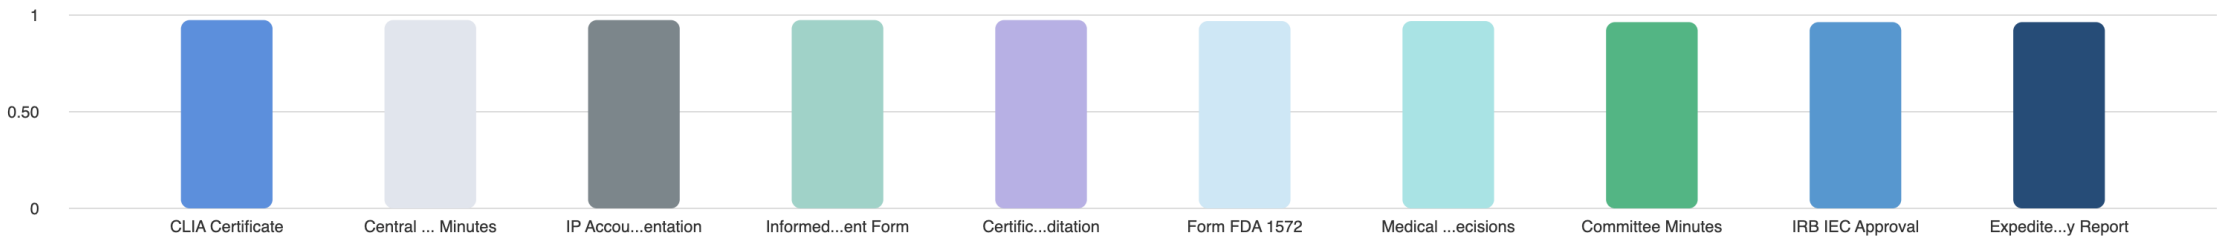
F1 Score Trends

Last Day

No trends data available

Document Performance Distribution

Bottom 10 | Top 10 | See All



Latest Validation Runs

See All

Run ID	Date	Status	Library Version	P95	Macro F1	Lowest Class F1
f5bb783a-1972-4...	05/05/2026		v1.41	97.4%	95.7%	94.7%
5f0bee5f-5a88-4...	05/05/2026		v1.41	97.4%	95.6%	94.7%
0060d99d-6c6a-4...	05/05/2026		v1.41	97.4%	95.6%	94.7%
e0f7efbf-f32b-4...	05/05/2026		v1.41	97.4%	95.5%	94.7%

How to Stay Validated?

Table 1: Control design stages.				
Stage 1	**Stage 2**	**Stage 3**	**Stage 4**	**Stage 5**
The system is used in parallel to the normal GxP processes	The system is executing a GxP process automatically but must be actively approved by the operator	The system is executing the process automatically but can be revised by the operator	The system is running automatically and controls itself	The system is running automatically and corrects itself

Table 2: Autonomy stages.					
Stage 0	**Stage 1**	**Stage 2**	**Stage 3**	**Stage 4**	**Stage 5**
Fixed algorithms are used (No machine learning)	The system is used in a locked state. Updates are performed by manual retraining with new training data sets	Updates are performed after indication by the system with a manual retraining	Updates are performed by automated retraining with a manual verifi cation step	The system is fully automated and learns independently with a quantifi able optimization goal	The system is fully automated and selfdetermines its task competency and strategy

The Verified Rule-Based Approach for TMF Completeness

Where We Are Now: Document-Level AI

- A probabilistic model classifies each document and extracts its metadata.
- The locked state, confidence thresholds and human review make that model inspectable.
- The Expected Document List is still a template, created during the study configuration.

Future: Study-Level Intelligence

- Expected Document List is computed from the live state of the study.
- A deterministic rule engine does the reasoning, and generative AI is used in the UI to help explain it.
- As the Autonomy increases, the Observability requirement rises alongside.
- A verified rule-based approach works for a neuro-symbolic system.

How Does a TMF Completeness Engine Work?

The TMF Plan & Protocol

These are processed and converted to a set of rules with HITL review.

Study Events

CTMS, safety, EDC and QMS feed events in that become triggers or placeholder fulfillments.

Rule Engine

- Triggers and cascade rules that create obligations.
- Each rule provides a citation showing the regulation or SOP.
- The database is deterministic, and the LLM simply narrates the result with citations.

Expected Document Lists

- EDLs are computed per study, country and site, and compared with what is filed.
- Select a date to see what was owed and what was known.
- The SOP that was effective and in force on that date applies.

Completeness Chat

- The AI answers based on the cascaded database of rules, obligations, evidence, and citations.
- Each citation provides document or regulation, version and effective date.
- The AI refuses to answer when the evidence is not there.

An Example…

A protocol amendment becomes effective, so the rule cascade raises IRB re-approvals, re-consent on ICF 2.0, an updated protocol signature page and an updated pharmacy manual. The AI can explain why the PI must be retrained on the updated protocol.

cdisc®

What an TMF Completeness Engine Can Do

Inspection Readiness

Reconstruct what the TMF should have contained on a given date, and what was known then. That is the form an inspector's question takes.

Mock Inspections

Where the regulation requires one thing and the plan or an SOP commits to another, the disagreement surfaces as a finding, with both citations.

SOP and Regulation Impacts

When an SOP is revised or a guideline changes, the affected obligations are listed automatically, by study, from the effective date.

Risk-Based QC Scope

QC effort follows artifact criticality and study risk rather than a flat sample, and the scope adjusts as signals change.

Archival Readiness

An inspection readiness determination showing the evidence and citations, instead of a simple set of completeness reports.

cdisc®

Validating a TMF Completeness Engine

Intended use	Compliance intelligence over the eTMF.
Risk assessment	Artifact criticality designations are customer-determined, with best practices assisting the AI engine.
Rule-set governance	Replaces dataset governance. Plans and SOPs are converted to rules and citations and verified by humans.
Performance qualification	The Expected Document List logic is tested against acceptance criteria for correctness.
Human oversight	While the AI provides guidance, it does not make the final decision. Humans verify what goes into the ruleset and what comes out of the expert system.
Audit trail	Essential for transparency of decision-making.
Change control	Rules, citations, effective versions, models and prompts.
Periodic review	Regulatory currency and rule updates feed the review cycle.

Approach

Treated as a neuro-symbolic expert system, where humans verify the rule sets going in.

The system combines the rule set with operational data from the eTMF and CTMS.

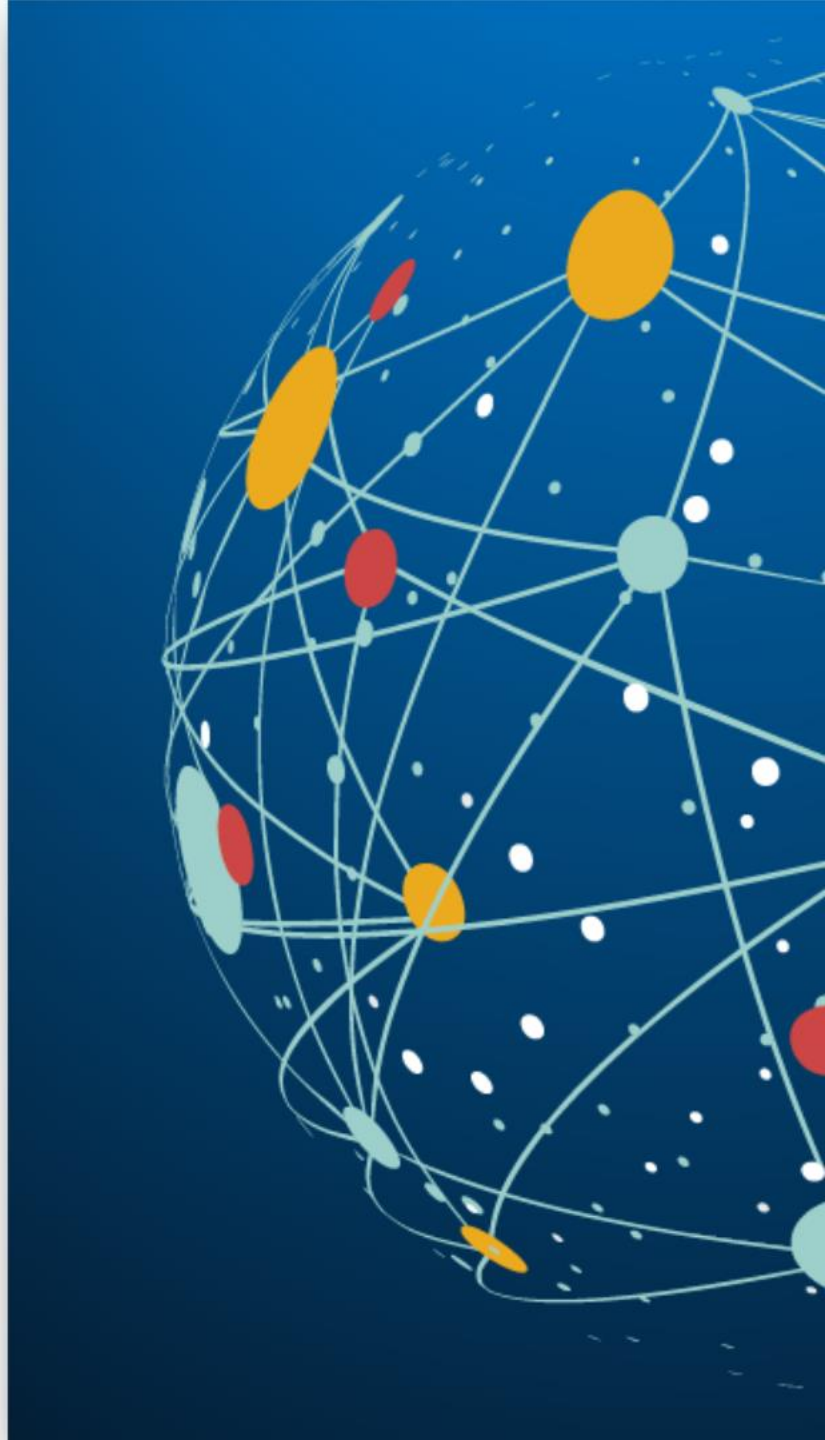
The AI simply narrates the resulting decision path to provide expert guidance.

UATs and other validations may be run against this deterministic information.

cdisc®

Actions and Next Steps

- There is not a new kind of validation. GAMP 5, Annex 11 and the ten principles still apply, but the approach must be different.
- For uncontrolled inputs, a locked state is the reason a probabilistic model can be validated at all. If the model, prompts, thresholds and mappings are fixed, it behaves like deterministic software.
- For controlled inputs, it is essential that grounding, citations and refusal to speculate are essential controls to make AI usable in GxP. It becomes more challenging when an LLM has uncertainty, conflicting rules, or runs at a temperature that is not ’cold’.
- Anything that reasons over your clinical and TMF data must build a rule database and not simply infer information; there can be zero creativity in responses.



Thank You!

