

DIGITALTRADE4.EU, RIHO VEDLER

Proposal and Comment Information

Title: Request for Comment on Model Changes for the Board's 2027 Supervisory Stress Test, OP-1882

Comment ID: FR-2026-0037-01-C04

Submitter Information

Organization Name: DigitalTrade4.EU

Organization Type: Foreign

Name: Riho Vedler

Submitted Date: 10/19/2026

DigitalTrade4.EU submits this comment in response to Docket OP-1882 concerning the Federal Reserve Board's proposed changes to the 2027 supervisory stress test models and FR Y-14 reporting requirements.

We support the objective of improving supervisory risk assessment while maintaining data quality, transparency and proportionate reporting obligations. Our submission focuses on how transaction driven activity measures and narrowly defined trade finance exposures can be identified consistently and reliably across financial institutions.

We recommend that data readiness should precede additional model complexity. Before introducing new reporting categories or activity measures, the Federal Reserve should evaluate the consistency of definitions, identifier coverage, transaction relationships, historical stability and comparability across institutions.

Our proposed Verified Trade Events framework distinguishes authoritative source records, verified commercial events or states, and the financial exposures or revenue measures associated with them. This approach supports reliable data classification and traceability without requiring the duplication of complete commercial documents. Existing identifiers, including LEI where appropriate, should be reused, while technologies such as vLEI should remain optional.

We propose a limited research pilot involving a small group of relevant financial institutions over two to four quarterly reporting cycles. The pilot would assess data quality, identification accuracy, potential improvements in supervisory models and actual reporting costs before any broader regulatory changes are considered.

Our recommendations do not seek preferential capital treatment for trade finance, mandatory adoption of new identity technologies or immediate additional reporting obligations. Instead, we advocate an evidence based and technology neutral approach that improves supervisory information quality while minimizing unnecessary complexity and compliance burden.

The complete submission provides a detailed implementation framework, proposed pilot data fields, evaluation criteria, confidentiality safeguards and practical recommendations for future supervisory model development.

Data Architecture for Future Stress-Test Model Development: Trade-Finance Exposure Identification and Transaction- Driven Activity Measures

DigitalTrade4.EU

Contact: Riho Vedler | secretariat@digitaltrade4.eu | Docket No. OP-1882 | October 2026

Abstract. DigitalTrade4.EU supports the Federal Reserve Board’s objective of improving the supervisory stress test’s ability to capture business-model heterogeneity while maintaining data quality, transparency and proportional reporting burden. This comment addresses a narrow implementation question raised by Docket OP-1882: how transaction-driven activity measures and narrowly defined exposures can be made reproducible across firms before their use in supervisory models. Based on DigitalTrade4.EU research into the data and interoperability conditions required to make commercially relevant trade events consistently identifiable and usable in regulated financial processes, we recommend a staged data-readiness approach: define the economic object, establish stable identifiers and relationship lineage, measure coverage and comparability, run a bounded model-development pilot, quantify both model utility and reporting burden, and only then consider model segmentation. For purposes of this comment, we refer to the underlying analytical framework as Verified Trade Events (VTE). The VTE framework distinguishes authoritative source data, a purpose-specific verified event or state, and any downstream financial decision, instruction or exposure to which that evidence may be relevant. A VTE is a purpose-specific statement that a defined commercially relevant event has occurred, or that a defined documentary, regulatory or commercial state currently holds, supported by evidence sufficient to resolve the authoritative source, the identity and authority of the relevant actor where applicable, record integrity and provenance, current status and transaction binding. It does not replace the underlying source record or determine a bank’s credit, compliance, legal or operational judgment. We add a simple measurement-error rationale for why inconsistent classification can obscure true revenue or risk relationships, propose an illustrative two-to-four-cycle pilot architecture and minimum field set, and recommend an extensible identifier-field design that can reuse established U.S. and international identifiers without mandating a single identity scheme or credential technology. We do not recommend favorable capital treatment for trade finance, a new stress-test factor, mandatory LEI/vLEI adoption, or Federal Reserve validation of underlying commercial events. The implementation annexes and pilot safeguards below add optional quantitative sensitivity settings, explicit cohort eligibility and zero-activity handling, frozen model-comparison baselines, historical and cross-border feasibility tests, confidentiality and data-stewardship controls, and a cost/decision-value ledger. None is offered as an empirically validated supervisory threshold or an expansion of the Board’s legal reporting perimeter.

KEY WORDS. stress testing; FR Y-14; noninterest income; trade finance; data architecture; measurement error; business-model heterogeneity; Verified Trade Events; entity identifiers; verifiable credentials; LEI; vLEI; interoperability; transaction-driven revenue

Summary of Principal Contributions

1. **Data readiness should precede model segmentation.** A narrow supervisory category should be used only after its semantics, identifier coverage, relationship mapping, missingness, current-state integrity and cross-firm comparability have been demonstrated.
2. **Transaction-driven revenue requires a common activity unit.** Before transaction count, processed value, documentary volume or a similar measure is used in a PPNR relationship, the service perimeter, unit, timing convention, corrections and cancellations must be defined consistently across firms.
3. **Trade finance is a pilot question before it is a model question.** The Board has identified that narrow lending types including trade finance cannot currently be identified reliably in the vendor estimation data. The appropriate next step is a bounded identification and data-quality pilot, not immediate model segmentation.
4. **VTE can support classification without document replication.** The relevant distinction is source record $\neq$ purpose-limited verified state $\neq$ financial exposure or revenue measure. A pilot can preserve references, current status and transaction binding without collecting complete trade documents.
5. **Identity should be an extensible field specification, not a new platform mandate.** The pilot can separate identifier scheme, identifier value/reference and verification method. LEI can support legal-entity matching where available; other authoritative schemes can coexist, while vLEI or other governed credentials are optional evidence for role or authority rather than mandatory FR Y-14 technology.
6. **Measurement quality is part of model identification.** Inconsistent classification can attenuate or otherwise distort estimated economic relationships. The pilot should therefore measure semantic consistency, deterministic linkage, unresolved records and historical stability before a variable is used in model estimation.
7. **A bounded pilot should produce an explicit go / revise / no-go decision.** A small, time-limited model-development exercise should measure supervisory utility and reporting burden in parallel. A positive result would support consideration through the Board's ordinary reporting and model-governance processes; a negative result would be evidence against broader collection. Pilot eligibility, diagnostic sensitivity bands, baseline comparisons, confidentiality controls and measured implementation costs should be documented before launch; none is a proposed universal reporting mandate or model threshold.

1. Scope and Position of the Comment

Docket OP-1882 invites public input on proposed changes to the Board's 2027 supervisory stress-test models and on targeted revisions to the FR Y-14A/Q/M reports that would facilitate future model development and improve risk capture. The proposal is particularly relevant to two issues addressed here: the move toward a more granular noninterest-income model that uses activity-volume information where economically appropriate, and targeted reporting changes intended to test whether additional distinctions should be incorporated into future models.

DigitalTrade4.EU's contribution is deliberately limited to data architecture. We do not offer a view on the calibration of the proposed noninterest-income equations, the mortgage-servicing-rights fields, or whether the proposed Facility Internal Risk Rating should represent loss severity or blended expected loss. We address the prior question that determines whether additional

granularity is reliable: can the relevant exposure, activity, counterparty or transaction state be identified consistently across firms and across time?

Our position is that a supervisory model should not use a narrow category merely because the category is economically plausible. A category becomes model-ready only after its semantics, identifier coverage, relationship mapping, current-state integrity, historical stability and incremental supervisory utility are demonstrated. This is especially important for trade finance, where a shipment, invoice, receivable, guarantee, documentary-control state, fee-generating service and financing exposure may relate to the same commercial relationship while remaining legally and analytically distinct objects.

2. Direct Relevance to the 2027 Model and Reporting Proposal

The 2027 package already follows a staged logic in several areas. For noninterest income, the Board proposes to replace the interim discount-factor approach with a more granular framework using volume-based or structural methods where the underlying business mechanism and data support them, while using alternative methods where suitable activity data are unavailable. The Board is also proposing Schedule G changes that may inform future PPNR model development. This creates a direct data-design question: what counts as a stable unit of transaction-driven activity, and can that unit be reported consistently enough to support cross-firm modeling?

For a transaction-driven service, a candidate volume measure has value only if it represents the same economic activity across reporters. “Transaction count,” “processed value,” “guarantee volume,” or “documentary-processing volume” may each be relevant in a particular business, but the unit, service perimeter, timing convention and treatment of cancellations, corrections and bundled services must be controlled before the measure can support a supervisory relationship with revenue.

The wholesale proposal provides the same lesson from the exposure side. The Board proposes new Schedule H.1 fields for Guarantee Amount, Securitization Flag and Facility Internal Risk Rating and seeks feedback on related mapping questions. The purpose is not to assume a model effect in advance, but to collect information that can test whether future model changes are justified. The Board’s review of prior comments also shows why this sequencing matters: missing or inconsistent data can prevent reliable estimation even where a risk distinction is conceptually plausible.

Schedule L provides a further data-governance precedent. The 2027 package addresses inconsistent reporting of client-cleared derivatives and inter-affiliate transactions by clarifying the population and relationship logic rather than introducing a new economic model solely to compensate for inconsistent reporting. This same discipline is useful for future trade-finance or transaction-service distinctions: first make the population and relationships reproducible, then test economic relevance.

Table 1. Mapping of DigitalTrade4.EU recommendations to the 2027 proposal.

2027 proposal area	Data problem	DigitalTrade4.EU contribution
Noninterest-income model	Transaction-driven revenues require activity measures that are economically linked to fee generation and comparable across firms.	Define the activity unit, service perimeter, stable linkage, time basis and mapping rules before relying on a volume measure in a model.
FR Y-14Q Schedule G	New information may be collected to support future PPNR development.	Use targeted machine-readable fields; measure classification quality and model utility before broadening the reporting perimeter.
FR Y-14Q Schedule H	Guarantees, securitization status and facility-level ratings require reliable exposure and relationship mapping.	Use stable exposure/entity identifiers and explicit relationship status so the same structure can be reproduced across firms.
FR Y-14Q Schedule L	Counterparty populations and transaction roles can be inconsistently classified across sub-schedules.	Treat identity and relationship lineage as a data-quality layer; record unresolved mappings rather than forcing false precision.

3. Trade Finance as a Data-Identification Case Study

The Board's model review identifies a measurement constraint directly relevant to this comment. In considering more granular corporate lending categories, the Board observed that narrower lending types such as trade finance cannot be identified in the vendor estimation data set and generally represent a small share of aggregate corporate exposure outside particular firms. The Board also cautioned that narrow categories can reduce model stability. We do not ask the Board to reverse that modeling judgment. We address the step before it: how future reporting could make a trade-finance subset consistently identifiable and therefore empirically testable.

A useful trade-finance identifier cannot be only a free-text product label. The minimum reproducible description may need to distinguish the financing structure, the financial exposure, the relevant obligor or guarantor, and - only where economically relevant - a bounded commercial or documentary state linked to that exposure. The same commercial transaction can generate multiple records and events, so the reporting architecture must preserve the distinction between the financial object used in the stress test and the source-side evidence used to classify or validate that object.

The supervisory value of a pilot is therefore not conditional on assuming that trade finance will ultimately receive a separate model treatment. A pilot can answer a prior empirical question that the Board currently cannot resolve from the vendor data: whether a narrow category can be identified consistently enough to test whether it has incremental explanatory, predictive or monitoring value. A negative result is also useful because it supports not expanding the reporting or model perimeter.

4. VTE Data Architecture: From Source Evidence to Model-Ready Attributes

For purposes of this comment, we refer to the underlying analytical framework as Verified Trade Events (VTE). The VTE framework separates authoritative source data, a purpose-

specific verified event or state, and the financial decision, instruction or exposure to which that evidence may subsequently be relevant. A VTE is a purpose-specific statement that a defined commercially relevant event has occurred, or that a defined documentary, regulatory or commercial state currently holds. It is not a replacement for the underlying document or authoritative source record, and it does not itself determine a bank's credit, compliance, legal or operational judgment. Rather, it is a bounded evidentiary statement supported by evidence sufficient to resolve the authoritative source; the identity and authority of the relevant actor, where applicable; record integrity and provenance; the relevant timestamp and version; current status, including correction or revocation where applicable; and binding to the transaction or relying context.

source data $\neq$ verified event / state $\neq$ financial exposure, revenue measure or instruction

This distinction is useful for supervisory reporting because it avoids two opposite errors. First, it avoids treating broad product labels as if they carried enough information to define a consistent model population. Second, it avoids solving the problem by copying complete trade documents into regulatory reporting. For many model-development questions, the Board may need only a small set of reproducible attributes: what financial object is being measured, what entity or relationship it involves, what bounded state is relevant, when that state was current, and whether the linkage can be resolved with sufficient confidence.

The VTE framework also treats entity resolution as an observable readiness variable rather than an assumption. Cross-system analysis should distinguish deterministic identifier coverage, the method used to link entities, match-confidence indicators and unresolved-identity incidence. This is directly relevant to any pilot that links guarantees, counterparties, financing structures or transaction-volume measures across systems.

5. Why Data Readiness Has Model Value: Measurement-Error Intuition

The case for a data pilot does not require an assumption that trade finance is lower risk. It rests on a more basic statistical point: if the activity or exposure used to estimate a relationship is measured inconsistently, the model may be unable to distinguish economic heterogeneity from reporting noise. For a transaction-driven noninterest-income component, consider the purely illustrative relationship:

$$\text{Revenue}_{i,t} = \alpha_i + \beta \cdot \text{Activity}_{i,t} + \gamma \cdot \text{Controls}_{i,t} + \varepsilon_{i,t}$$


If the reported activity measure is a noisy proxy for the true activity, so that $\text{ReportedActivity}_{i,t} = \text{Activity}_{i,t} + u_{i,t}$, the standard classical measurement-error result is:

$$\text{plim}(\hat{\beta}) = \frac{\beta \cdot \text{Var}(\text{Activity})}{\text{Var}(\text{Activity}) + \text{Var}(u)}$$

Here $\hat{\beta}$ denotes the estimated activity coefficient and $\text{plim}(\hat{\beta})$ denotes its probability limit as the sample size grows. $\text{Var}(\text{Activity})$ denotes the variance of the latent true activity measure over the observations that supply identifying variation in this illustrative regression, while $\text{Var}(u)$ denotes the variance of the measurement error under the classical assumptions. In an actual panel specification, the relevant identifying variation depends on the model

transformation, controls and fixed effects; the expression is therefore a pedagogical simplification rather than a proposed supervisory equation.

Plain-language interpretation. If firms classify the same economic activity differently, or combine economically different activities under the same reported measure, the input contains noise. Even when true activity is related to revenue, that noise can make the estimated relationship appear weaker. The model may then fail to detect real business-model heterogeneity because the input is not comparable, rather than because the underlying economic relationship is absent.



Implication for OP-1882. This mechanism explains why a new transaction-count, processed-value or other activity-volume field should be tested for semantic consistency, missingness, deterministic linkage, correction handling and historical stability before it is relied upon in a transaction-driven revenue model. Poor input consistency can attenuate or destabilize an estimated relationship; a data-readiness pilot therefore tests whether the field is capable of identifying economic heterogeneity before the Board uses it to justify additional model segmentation.

Real-world classification error may be non-classical and can bias estimates in directions other than attenuation. The point is not that this formula describes every FR Y-14 error. The narrower point is that classification quality is part of model identification, not merely a post-model data-cleaning issue. A pilot can measure the magnitude and structure of that error before the Board decides whether a narrow category or activity measure belongs in a supervisory model.

Non-classical error diagnostic. The pilot should not infer that improved coding necessarily produces classical attenuation. Missingness may depend on firm size, technology, product mix or stressed economic conditions; product misclassification may be asymmetric; and duplicate or reversed records may affect both the regressor and outcome. Where a sufficiently reliable adjudicated sample exists, compare reported classes with source-backed classes in a confusion matrix, test reporter- and period-specific error rates, and examine sensitivity to plausible differential misclassification and unresolved links. A latent-class or mixture specification could be considered only if sample size, independent information and identification assumptions permit it; it is not a prerequisite for the bounded pilot.

Supervisory utility test

A pilot is useful if it shows whether a proposed field or category materially improves model-development or monitoring capability relative to a simpler baseline - for example, by improving out-of-sample fit, reducing unexplained cross-firm dispersion attributable to reporting differences, improving population stability, or enabling a previously untestable hypothesis to be evaluated.

The Board need not assume a capital effect in advance. If the proposed distinction does not improve risk or revenue capture, the pilot itself provides evidence against broader collection.

For evaluation, freeze a comparison design before examining results: B0 is the current feasible coarse-category or activity-volume baseline; B1 applies harmonized definitions and quality controls while leaving the same baseline predictors; B2 adds the candidate narrow trade-finance or service-volume attributes to B1. Comparing B1 with B0 isolates the contribution of data harmonization; comparing B2 with B1 isolates incremental candidate-field value. Where histories permit, compare errors on held-out periods and test stability across reporters, using appropriate revenue-model prediction errors or exposure-risk diagnostics rather than R^2 alone. With only three to five firms and a few quarterly cycles, these analyses are feasibility diagnostics, not validation of an across-cycle supervisory stress model.

6. Organizational Identity: LEI and a Limited Role for vLEI

Where legal-entity identity is relevant, the pilot should use an extensible identifier field rather than create a new trade-finance-specific identifier or hard-code one universal scheme. The reporting object can separate *Entity_ID_Scheme* from *Entity_ID_Value_or_Reference* and, where useful, the authoritative registry or issuer and an as-of/status field. LEI is a strong option where it is available. Other identifiers may be relevant only for defined populations - for example, SEC CIK for EDGAR filers, SAM.gov Unique Entity ID (UEI) or CAGE in federal award/procurement contexts, or jurisdictional business-registration identifiers. The pilot should not assume that any one of these has universal coverage, and sensitive tax identifiers should not be treated as a default cross-system key.

A different issue arises when the data-quality question concerns not only which legal entity is involved, but whether a person or service is acting with a defined organizational role or authority. GLEIF's verifiable LEI (vLEI) ecosystem provides governed, machine-verifiable credentials for legal entities and organizational roles anchored to the LEI. That capability is conceptually aligned with the VTE requirement to resolve identity and authority, but OP-1882 does not require the Board to create a role-credential regime. A future pilot could therefore record a verification method or evidence reference where role or authority is material, while keeping vLEI optional and avoiding any general FR Y-14 credential mandate.

Generic verifiable-credential standards and vLEI should also not be conflated. The W3C Verifiable Credentials Data Model is a general framework for issuer-signed, machine-verifiable claims; W3C expressly distinguishes cryptographic verification of a credential from validation of whether its claims are true and suitable for a verifier's purpose. vLEI is a separate governed organizational-identity ecosystem. Accordingly, a credential can carry a signed claim that two identifiers, an entity and a role, or an entity and an authority relationship are linked, but the reliability of that mapping still depends on whether the issuer is authoritative for the claim, whether the credential is current, and whether the verifier's reliance rules accept it.

For OP-1882, the practical implication is a field-level design rather than an infrastructure prescription. An illustrative pilot may use *Entity_ID_Scheme*, *Entity_ID_Value_or_Reference*,

Identity_Verification_Method, Verification_Status_or_AsOf and - only where role or authority is material - *Role_Authority_Evidence_Reference*. The Board would not need the complete credential or source-registry record merely to test classification quality; the pilot should capture only the information needed to reproduce the mapping and assess its reliability, subject to applicable confidentiality, access and recordkeeping requirements.

7. Recommendations

Recommendation 1 — Apply a data-readiness gate before adding granular model segmentation. Before a new exposure or business-line category is used in a supervisory model, the Board should assess semantic consistency, coverage, missingness, deterministic mapping rate, unresolved-link incidence, temporal stability and cross-firm comparability. These checks should be treated as prerequisites to model complexity rather than as post-model data cleaning.

Recommendation 2 — Tie transaction-driven revenue measures to the economic unit that generates the fee. The proposed noninterest-income framework appropriately moves toward direct activity measures where possible. For transaction-driven services, a candidate volume measure should have a clear economic relationship to the revenue line, a stable unit of observation, a governed definition across firms and a reproducible time basis. Transaction count, processed value, guarantee or documentary-processing volume may be suitable in particular businesses, but none should be presumed comparable without a data-quality test.

Recommendation 3 — Use a bounded research or model-development pilot before any mandatory trade-finance segmentation. The Board could test a minimum field set with a small cohort of firms that have relevant trade-finance or transaction-services activity and a defined period of two to four quarterly cycles, supplemented by historical backfill where feasible. The exact cohort and duration should be set by the Board to capture sufficient business-model heterogeneity. The pilot should test data readiness, model-development utility and reporting burden; it should not be treated as a new risk-weighting category or as sufficient by itself to validate a long-horizon stress model.

Recommendation 4 — Use an extensible identity field specification and reuse authoritative identifiers rather than mandate a new identity system. A pilot should first use existing facility, exposure and counterparty keys maintained by firms. Where legal-entity identity is needed, report the identifier scheme and value/reference separately so LEI and other context-appropriate authoritative schemes can coexist. A verification-method/status field can record whether the link was established through a direct registry match, a controlled reporting-firm mapping, or an optional governed credential. vLEI may be evaluated where organizational role or authority is relevant, without becoming a general FR Y-14 requirement.

Recommendation 5 — Preserve current-state and relationship lineage. If a field depends on an external commercial or documentary state, reporting should identify the relevant state type, effective timestamp, current status/version, source class and transaction/exposure linkage. The reporting process should make unresolved or confidence-limited links visible rather than silently imputing certainty.

Recommendation 6 — Follow a define → measure → pilot → model sequence. A new data distinction should move to formal stress-test treatment only after it demonstrates practical utility, sufficient historical depth and a measurable improvement in risk or revenue capture that justifies added complexity and reporting burden.

Recommendation 7 — Measure burden inside the pilot. Rather than estimate compliance cost abstractly, the Board could require pilot participants to report one-time implementation effort, recurring preparation time, manual intervention, source-system count and the share of

fields already available in existing reporting or authoritative systems. This would turn PRA proportionality into an observed pilot outcome.

Table 2. Illustrative data-readiness gate for a future narrow exposure or activity category.

Gate	Quantitative / observable metric	Pilot decision rule
Semantic consistency	Share of eligible observations mapped to a controlled definition; exception rate; reporter-specific overrides.	Assess eligible-observation mapping and exceptions at both reporter and pooled levels; run a pre-specified range of diagnostic settings (Table 2A). Do not accept incompatible economic definitions just because pooled coverage is high.
Identifier coverage	Deterministic link rate = deterministically linked eligible records / eligible records; unresolved-link rate = unresolved eligible records / eligible records.	Compare deterministic and unresolved rates against Table 2A sensitivity settings by reporting firm, time and product. Escalate concentrated gaps; no numerical setting is a universal acceptance rule.
Current-state integrity	Timestamp/status coverage; stale, superseded or corrected-state incidence.	Check as-of reproducibility and correction/revocation sequencing; pause use where stale states cannot be isolated for the research question.
Cross-firm comparability	Missingness, mapping exceptions and outlier rates by reporter; dispersion before/after harmonized mapping.	Report reporter-level distributions and eligibility denominators; redo mapping or reject a field if the distinction mainly reflects divergent reporting rules.
Temporal stability	Definition and mapping breaks across reporting cycles; historical backfill coverage.	Record each schema/ownership break and compare available-history cohorts. Do not extrapolate a short clean backfill into a full stress horizon.
Incremental model utility	Change in fit, stability, validation performance or monitoring value versus a simpler baseline.	Use pre-specified B0/B1/B2 comparisons, withheld observations where feasible, and stability diagnostics; distinguish data-cleaning gains from new-field gains.

Gate	Quantitative / observable metric	Pilot decision rule
Burden-to-utility	One-time implementation hours; recurring preparation hours; manual-touch rate; source-system count; share reused from existing data.	Compare observed full pilot cost, including Board and vendor time where estimable, with decision-relevant information gain; stop or narrow the pilot if cost is disproportionate.

Table 2A. Optional sensitivity settings for pilot diagnostics (illustrative only; not regulatory cutoffs).

Diagnostic	Illustrative values to test	Interpretation and safeguard
Deterministic eligible-record linkage	90% / 95% / 99%	Compare model and quality conclusions at each setting; report unmapped denominator and disaggregate by firm.
Unresolved eligible-records	10% / 5% / 1%	Test how selective exclusion or explicit unresolved coding changes conclusions; unresolved is not zero exposure.
Semantically harmonized eligible records	90% / 95% / 99%	A high percentage cannot override material definition mismatches, especially for an individual participant.
Status and as-of reconstructability	90% / 95% / 99%	Use only for state-dependent fields; exclude fields whose authoritative status cannot be determined at the relevant date.

These values are deliberately round sensitivity points, not empirically calibrated supervisory standards or DigitalTrade4.EU estimates of current industry performance. Before any pilot decision, the Board would determine applicability, denominators, tolerated exceptions and statistical uncertainty for each use case. No single percentage overrides legal validity, material data gaps, sample selection or actual model utility.

8. Limited Research Pilot: Implementation Architecture and Candidate Fields

Pilot scope. The pilot should be bounded and explicitly framed for data-quality and model-development research rather than as an immediate recurring FR Y-14 obligation. A small cohort could, illustratively, include several firms with material trade-finance or transaction-services activity; three to five participants may be a workable starting point, but the Board should choose the cohort based on heterogeneity and data availability rather than a fixed number. Two to four quarterly reporting cycles can test definition stability, mapping quality and recurring burden. That period is not, by itself, enough to establish stress-model performance across economic cycles, so historical backfill should be used where feasible and appropriate. Eligible respondents should be firms within the applicable FR Y-14 reporting and supervisory perimeter that actually conduct material relevant activity; a firm with no eligible activity should not be compelled by this proposal to build trade-finance reporting infrastructure. The design should distinguish confirmed zero eligible observations, non-applicability, non-response and unavailable data rather than collapsing them into a missing value. A small purposive sample is not statistically representative of all FR Y-14 filers.

Collection vehicle and intended use. The Board could use a separate pilot template, controlled data call or other legally appropriate model-development mechanism before proposing changes to recurring FR Y-14 schedules. The collection method should be selected consistently with the Board’s authority and applicable Paperwork Reduction Act requirements. Scope, data use, confidentiality treatment and decision criteria should be stated in advance. A positive pilot would not automatically move a field into FR Y-14 or a supervisory model; it would support a later decision whether to propose, revise or reject a broader collection through the ordinary process. Pilot submissions should be managed through, or explicitly reconciled with, existing authorized FR Y-14 data-governance controls, including definitions, edits, reconciliation, validation exceptions, permitted access, change logs and disposition of data at the end of its authorized retention period. An exploratory template must not silently become a new recurring regulatory obligation.

Table 3. Illustrative bounded research-pilot implementation architecture.

Stage	Core task	Output / decision evidence
1. Taxonomy and mapping	Define the eligible population, service perimeter, activity unit, identifier schemes, state definitions, timing conventions, source mappings and known exceptions. An initial design period of roughly one quarter may be sufficient for this step, subject to participant complexity. Publish the inclusion/exclusion rules, reporter-selection rationale, optional cross-border test case, fixed B0/B1/B2 baseline specification, diagnostic bands, and signed-off data-protection plan.	Controlled data dictionary; source-to-field map; unresolved definitions; initial burden baseline.
2. Limited collection	Collect the minimum candidate fields for two to four quarterly cycles through the selected pilot mechanism. Permit historical backfill where existing source data can be mapped without disproportionate reconstruction. Distinguish participants without eligible trades from firms with missing data; track retrospective reconstruction by lookback cohort and any manual/vendor workaround.	Pilot dataset; linkage and exception logs; definition-change history; observed recurring preparation effort.

Stage	Core task	Output / decision evidence
3. Readiness, utility and burden assessment	Apply the Table 2 data-readiness metrics, cross-firm comparability checks and the Section 5 measurement-error logic. Test whether the fields improve monitoring or model-development utility relative to a simpler baseline, while measuring burden using the later burden metrics. Use the B0/B1/B2 design, non-classical error sensitivities, eligible-record denominators and confidential reporter-level diagnostics; add pilot-cost and information-value scorecards.	Data-quality scorecard; model-utility comparison; burden evidence; identified failure modes and required definition changes.
4. Decision gate	Choose among no-go, revise-and-retest, or consideration of a formal reporting/model proposal. Any broader FR Y-14 or model use would remain subject to the Board's normal legal, PRA, validation and public-input processes.	Documented decision explaining whether incremental supervisory utility is proportionate to data quality and burden. Record the controlled access/retention plan, remaining legal or data gaps, and why broader collection would or would not merit a subsequent public process.

This sequence is intentionally reversible. The pilot is successful if it produces a defensible decision, including a no-go decision where the category cannot be identified consistently, the field adds little incremental model value, or the burden is disproportionate.

Candidate field set. The following fields are intended as a concrete pilot proposal, not as a request to add them immediately to the current FR Y-14 schedules. They illustrate how the Board could test whether a trade-finance population or transaction-driven activity measure is reproducible without collecting complete commercial documents. Where an equivalent field already exists in FR Y-14 or a firm's source systems, it should be reused rather than duplicated.

Before creating any new field, the Board should also inventory potentially reusable information in existing prudential reports, including FR Y-14, Call Report and other relevant collections. Reuse should occur only where the reporting perimeter, definition, unit and as-of date are demonstrably aligned. The same discipline applies to identity data: an existing identifier should be reused only where its population, authority and permissible use fit the pilot. This comment does not propose a field-by-field FFIEC 009 or Call Report crosswalk because equivalence must be established empirically rather than assumed from similar labels.

Table 4. Candidate fields for a limited trade-finance / transaction-activity pilot.

Candidate field	Possible schedule link	Purpose / boundary
Existing facility / exposure identifier	Schedule H or existing firm key	Stable financial-object link. Reuse existing keys; do not create a new universal identifier unless necessary.
Controlled financing / service structure code	Schedule H research field; Schedule G activity taxonomy	Distinguish defined structures using controlled enumerations rather than free text.
Entity identifier scheme + value / reference	Schedule H; relevant Schedule L research linkage	Separate scheme from value/reference so multiple authoritative identifier systems can coexist. LEI may be used where available; other schemes should be accepted only for their defined population. Do not treat a credential format as the identifier itself.
Guarantee amount and binding status	Schedule H.1 Guarantee Amount plus pilot linkage status	Connect the reported amount to the identified guarantor and exposure; avoid inferring legal effectiveness from amount alone.
Relevant commercial / documentary state type	Pilot only, where economically relevant	Examples: invoice accepted, delivery confirmed, documentary control state. Preserve legal distinctions; do not treat all as one “trade event.”
State timestamp, version and current-status flag	Pilot only	Supports temporal ordering and prevents stale or superseded states from being treated as current.
Identity / source verification method and relationship confidence	Pilot metadata	Record how an identity or relationship was established (for example, direct authoritative-registry match, controlled firm mapping or governed credential/attestation), its current/as-of status, and unresolved/confidence-limited links. A cryptographically verifiable credential does not by itself establish the truth of the asserted mapping.
Transaction-driven activity measure	Schedule G research field	Examples may include transaction count, processed value or service-specific volume only where a stable link to the fee-generating activity is established.

8.1. Participant selection, eligible population and proportionality

Sampling design. If authorized and operationally appropriate, select an initial three-to-five-firm purposive cohort using a written matrix of relevant-activity materiality, business-model and product mix, data-readiness maturity, and U.S.-reporting perimeter. Include different organizational types only where a valid reporting or participation basis exists; foreign banking organizations may be represented through the relevant U.S. reporting entity where applicable, not by presumed overseas reporting duties. Explain any absence of a category rather than implying representativeness. Participation terms, including whether voluntary or otherwise authorized, should be resolved under the Board’s applicable procedures.

Eligibility and zero activity. Define the eligible universe using a reporting-entity × product/service × period rule. Count (i) genuinely no eligible trade-finance exposure or activity, (ii) an in-scope activity with an observed value of zero, (iii) missing/unavailable observations and (iv) unresolved classification separately. Do not impose placeholder reporting on firms outside the pilot or treat a non-bank trading platform as a FR Y-14 respondent. Where third-

party or non-bank providers supply source evidence, the reporting bank remains responsible for its data lineage and can reference that evidence where legally permitted.

8.2. Candidate controlled vocabularies and entity life-cycle tests

Provisional codes, not a new financial-product standard. Keep two separate facets: a financing/contingent-exposure structure (for example, funded receivables or invoice financing; funded payables/supply-chain financing; letter-of-credit-related contingent commitment; independent guarantee or standby instrument; other; unknown/unresolved) and the fee-generating service/activity unit (for example, issued instrument, amendment, presentation/examination, payment processed, financing draw, outstanding value at as-of date). A single trade may generate more than one service event; a contingent guarantee is not automatically a funded loan; counts, balances, gross values and fee events cannot be combined into one denominator. Local product labels should map through documented crosswalks. ISO 20022 and ICC rules may inform specific fields or instrument conventions where applicable, but neither is represented as a universal supervisory taxonomy.

Entity life cycle. Test effective-dated links across mergers, acquisitions, spin-offs, legal name changes, domicile changes and successor identifiers. Preserve original and successor identifiers, valid-from/valid-to timestamps, source of the corporate-action mapping and unresolved legal relationships. Compare both the as-reported and historically restated populations; do not silently merge predecessor and successor exposures into a current identifier for backtesting. A few deliberately selected corporate-action cases can reveal whether reproducibility survives restructuring.

8.3. Historical reconstruction, cross-border evidence and technology readiness

Historical backfill. Start with an inventory of available history and retention restrictions, not a fixed retrospective collection mandate. Report the number of complete reporting periods available under the present definitions, separately for prospectively captured records and reconstructed records; a feasibility scorecard may distinguish 0–4, 5–12 and 13 or more available quarters as descriptive bins, not adequacy requirements. Record reconstruction source, methods, definition breaks, imputation, manual review effort and vendor limitations. If stressed-cycle observations are unavailable, label any model evaluation as limited to available conditions.

Cross-border comparability. Add a narrowly scoped test for relevant U.S.-reportable cross-border activity, including whether international identifiers, transaction states and time conventions map to the same U.S. reporting definition. Apply source-local verification, reference-based evidence and jurisdiction-specific confidentiality and transfer restrictions; avoid exporting original commercial documents or personal information merely to populate a pilot. Any GDPR or other foreign-law obligation depends on the actual parties and processing facts and is not assumed to govern every U.S. reporter. Existing international standards may be considered for interoperability, but this comment proposes no Basel/EU harmonization mandate.

Operational feasibility. Each participant should distinguish automatically generated, vendor-derived and manually reconciled fields, identifying minimum system changes, external dependencies, failure modes, and whether a low-technology extraction can produce the same semantics. A pilot can terminate an expensive field even if a few large firms can report it. Limit follow-up to the reporting perimeter in scope; heterogeneity in smaller or less complex entities should inform reporting proportionality, not create new obligations for non-reporters.

8.4. Confidentiality, stewardship and model-governance hand-off

Confidentiality by design. Collect purpose-limited derived values and controlled references where they answer the research question; use full source documents or party-identifying attributes only when justified by a specific authorized test. Maintain a data inventory and classification for commercially sensitive pricing, counterparty, shipment and relationship information; assign an accountable data steward, least-privilege access, audit logging, secure transfer, encryption, correction/version controls, and an authorized retention and deletion/disposition schedule. Pseudonymized tokens and hashes should not be assumed anonymous. The lawful confidentiality and disclosure treatment must be determined before collection, not inferred from technical safeguards.

Validation and decisions. Agree the pilot data dictionary and quality/edit rules with the existing supervisory data processes and identify which controls are reused versus genuinely pilot-specific. Retain versioned source-to-field maps, rejected-record and override logs, a frozen analysis plan, reviewer challenges and reproducible code or specifications under appropriate access restrictions. A no-go/revise/consideration decision should show reporter-level readiness, cross-border transfer issues where applicable, tested model utility, the full burden ledger and remaining limitations. Any resulting reporting or material model changes would require the Board's normal legal, PRA, model-review and applicable public-input steps.

9. Worked Example: Guarantee-Backed Invoice Finance

Consider an invoice-finance exposure for which a guarantee covers part of the facility and the lender's workflow also records an invoice-acceptance state. The supervisory object is the financial exposure, not the invoice or the commercial event. A reproducible pilot record could therefore anchor on the existing facility identifier, link the obligor and guarantor through stable entity identifiers, report the guarantee amount under the applicable Schedule H field, and carry only a bounded status reference to the invoice-acceptance state if that state is part of the research question.

The purpose of the example is not to suggest that invoice acceptance reduces probability of default or loss given default. It shows how a future study could determine whether exposures sharing a defined financing structure are consistently identifiable across firms, and whether any observed outcome difference remains after controlling for borrower, facility and other relevant characteristics. If the state or guarantee link cannot be resolved reliably, the record should be flagged as unresolved rather than included through a forced match. If a guarantor merges during the observation window, the example should retain the effective-dated predecessor/successor linkage rather than overwrite historical observations. A bank may retain the invoice in its authorized source environment while supplying a bounded, permissioned state reference to the pilot.

Table 5. Example data lineage for a pilot observation.

Layer	Example attribute	What the Board would use it for
Financial object	Facility / exposure ID; balance; internal risk rating	Defines the stress-test observation.
Entity relationship	Obligor / guarantor identifier scheme and value/reference; LEI where available; guarantee amount; verification method/status where relevant	Makes the guarantee relationship reproducible while separating the identity scheme from the method used to verify the mapping.

Layer	Example attribute	What the Board would use it for
Purpose-limited state	Invoice accepted; timestamp; current-status/version flag	Allows a defined research variable without collecting the full invoice.
Lineage / quality	Source class; linkage method; unresolved/confidence flag	Lets model developers distinguish missingness and mapping uncertainty from economic outcomes.

10. Paperwork Reduction Act Considerations

The proposed architecture aligns with the Board’s Paperwork Reduction Act questions. Practical utility is strongest when each field is tied to a defined model-development, monitoring or reconciliation purpose. Quality and clarity improve when definitions are controlled, identifiers are stable, status semantics are explicit and unresolved mappings are visible. Burden can be reduced when firms reuse data already maintained in authoritative internal or external systems and submit purpose-limited attributes instead of duplicating full source records.

Automation should follow semantic standardization rather than substitute for it. Machine extraction of inconsistent categories merely accelerates inconsistent reporting. A smaller number of well-defined, reusable fields can therefore have greater supervisory value than a larger collection of weakly standardized attributes.

A staged pilot also provides a burden-control mechanism. The Board can test coverage, mapping quality, historical depth and predictive or explanatory value with a bounded cohort and period before deciding whether any field belongs in a broader FR Y-14 collection. Pilot participants should record burden in parallel with data-quality and utility measures. A go decision should mean only that the evidence is sufficient to consider the next formal step; it should not predetermine a reporting amendment or model change. The pilot should measure burden directly rather than rely on speculative cost estimates.

Table 6. Illustrative pilot burden measures.

Burden measure	What to observe	Why it matters
One-time implementation effort	Hours to map definitions, connect source systems, build controls and validate lineage.	Separates setup cost from recurring reporting cost.
Recurring preparation effort	Hours per reporting cycle, including review and exception handling.	Supports PRA estimates based on observed effort.
Manual intervention	Share of eligible records requiring manual classification, identity resolution or relationship repair.	High manual-touch rates signal poor scalability and weak standardization.
Source-system footprint	Number and type of systems required to populate the pilot fields.	Shows whether a seemingly small field set creates disproportionate integration burden.
Reuse ratio	Share of pilot fields already available in FR Y-14, other prudential reporting, or authoritative internal systems.	Rewards reuse and helps identify genuinely incremental data requests.
Pilot administration and control effort	Board and respondent review, validation, legal, privacy/security and coordination hours, where measurable.	Includes work needed to run a legally and operationally sound pilot, not only populate fields.

Burden measure	What to observe	Why it matters
Vendor and one-off cash expenditure	Incremental software, vendor extraction, licensing, transfer or testing fees actually incurred.	Avoids assuming every cost appears as employee hours.
History and jurisdictional friction	Reconstruction hours and unsupported periods; extra cross-border review or permitted-access steps.	Makes backfill feasibility and conditional cross-border burden visible.

10.1. Cost and decision value: an evidence-based pilot ledger

Cost calculation. For each participant, estimate one-time effort as recorded hours by task and staff category multiplied by documented internal cost rates, plus incremental vendor or external charges; estimate recurring effort as measured hours and fees per cycle multiplied by the intended frequency. Report Board-side pilot design, validation, secure-handling and coordination effort separately where measurable. Present total, per-participant and per-eligible-observation figures, and sensitivity to the number of cycles and participants. Do not insert unverified industry-wide dollar estimates or generalize a three-to-five-firm sample to all FR Y-14 filers.

Value-of-information decision. The benefit of the pilot is the quality of the subsequent reporting/model decision, not an assumed dollar saving or capital reduction. State before launch the decision uncertainties to be resolved: whether identification is feasible; whether B1 outperforms B0 solely through better mapping; whether B2 adds stable value over B1; and whether any incremental value justifies the observed burden. Use pre-defined stop/revise signals for persistent incompatibility, inadequate permitted access, excessive manual handling or negligible incremental utility. A quantitative net-benefit estimate is warranted only if credible cost and benefit evidence later exists; otherwise report a transparent burden-versus-utility scorecard and its uncertainty.

11. Boundaries of the Recommendation

Nothing in this comment should be read as recommending lower loss severity, lower risk weights, favorable capital treatment or a separate stress-test category for trade finance. The VTE framework treats the magnitude and economic consequences of verification and financial-control frictions as empirical questions. Better data lineage may improve classification and reduce operational ambiguity; it does not establish lower borrower credit risk.

Nor should the Board become responsible for determining whether goods were delivered, whether an invoice is commercially valid, whether a negotiable document transfers legal rights, or whether a particular attestation is sufficient for a private financing decision. Those questions remain with the relevant parties, source systems, legal regimes and regulated institutions. The Board's interest for stress-test purposes is narrower: whether a reported attribute is defined, traceable, current and comparable enough for the stated supervisory purpose.

The recommendations are technology-neutral. LEI is cited as an established legal-entity identifier; other context-specific authoritative identifiers can coexist through a scheme-and-value field design. vLEI is cited only as an optional governed mechanism for organizational identity and role/authority evidence, while W3C Verifiable Credentials is a separate general standards family for machine-verifiable claims. Neither should be read as a required FR Y-14 stack. A credential authenticates an issuer's statement and status; it does not, by cryptography alone, prove the substantive truth of an identifier mapping or commercial event. Existing bank,

enterprise, payment, document and public-sector systems can remain authoritative for their respective functions.

The measurement-error equation in Section 5 is an explanatory device, not a proposed Federal Reserve model specification or a claim that all classification errors are classical. Its purpose is to show why comparability and classification quality must be tested before a narrow activity measure is treated as economically uninformative. Similarly, the data-readiness metrics in Table 2 are intended to support explicit pilot governance; numeric thresholds should be calibrated by the Board to the proposed use, materiality and validation standard rather than imported as universal cutoffs. Table 2A contains optional sensitivity settings rather than a supervisory threshold recommendation. The proposed controlled vocabularies, cohort design and cost ledger are technical options for assessing feasibility and do not create a duty for non-bank providers, non-participants or entities without eligible activity.

12. Conclusion

OP-1882 demonstrates the Board’s effort to make the supervisory stress test more transparent, granular and sensitive to differences in firms’ business activities. Greater granularity increases the importance of the classification and identity layer beneath the model. DigitalTrade4.EU recommends that the Board make data readiness an explicit gate: define the economic object, establish stable and current relationships, measure missingness and comparability, pilot the data, quantify both model utility and reporting burden, and only then decide whether additional model segmentation is warranted.

The VTE framework supports this recommendation by providing a precise separation between source records, purpose-limited verified states and downstream financial objects. That separation can help future FR Y-14 research identify trade-finance exposures or transaction-driven activity without forcing the Board to collect complete commercial documents or to validate commercial events. An extensible identity specification can reuse LEI or other authoritative identifiers appropriate to the reporting population, while optional role/authority credentials remain subordinate to the Board’s actual modeling purpose and to the authority of the issuer that makes the claim.

Accordingly, we recommend a bounded pilot rather than immediate mandatory reporting: define the taxonomy, select a small and sufficiently heterogeneous cohort, collect the minimum fields for a limited number of cycles, preserve state and relationship lineage, measure unresolved mappings and compliance burden, and compare incremental model or monitoring utility against a simpler baseline. The final gate should explicitly allow no-go, revise-and-retest, or consideration of a formal proposal. This makes the pilot useful even if the result is that no additional FR Y-14 collection or model segmentation is warranted. Documented cross-border applicability, explicit zero/non-applicable status, tested entity life-cycle mappings, pre-specified baseline comparisons, proportional confidentiality safeguards and actual pilot-cost accounting should accompany that decision. A non-bank source of supporting evidence does not thereby become an FR Y-14 reporting entity.

References

1. Board of Governors of the Federal Reserve System (2026a). “Request for Comment on Model Changes for the Board’s 2027 Supervisory Stress Test,” Docket No. OP-1882, 91 Fed. Reg. 62729–62733 (Oct. 2, 2026).
2. Board of Governors of the Federal Reserve System (2026b). “Dodd-Frank Act Stress Tests 2027,” model documentation and proposed model changes, Fall 2026.
3. Board of Governors of the Federal Reserve System (2026c). “Review of Comments and Summary of Changes to the Proposed Stress Test Models,” Fall 2026.

4. Board of Governors of the Federal Reserve System (2026d). “Enhanced Transparency and Public Accountability of the Supervisory Stress Test Models and Scenarios,” final rule and policy statements, Fall 2026.
5. Global Legal Entity Identifier Foundation (GLEIF) (2026). “The verifiable LEI (vLEI)” and vLEI Ecosystem Governance Framework v4.0.
6. World Wide Web Consortium (W3C) (2025). “Verifiable Credentials Data Model v2.0,” W3C Recommendation, 15 May 2025.
7. U.S. Securities and Exchange Commission (SEC) (2026). “CIK Lookup” and EDGAR filer-identifier guidance.
8. U.S. General Services Administration, SAM.gov (2026). “Entity Registration” and Unique Entity ID guidance.
9. Federal Acquisition Regulation, Acquisition.gov (2026). FAR 52.204-16, “Commercial and Government Entity Code Reporting.”
10. DigitalTrade4.EU (2026). “Verified Trade Events and Trade-Finance Interoperability: Working Capital and End-to-End Frictions in Global Value Chains,” working paper, version 14.