ARIADNE

ARIADNE DATAWORKS LTD.
CONTACT@USEARIADNE.COM

PUBLIC STATEMENT · CONSULTATION RESPONSE

Swap and Security-Based Swap Data Reporting

91 FR 37877 · CFTC RIN 3038-AF70 · SEC RIN 3235-AN78 · RELEASE NO. 34-105734 · FILE NO. S7-2026-22

July 6, 2026

VIA REGULATIONS.GOV
Christopher Kirkpatrick, Secretary of the Commission
Commodity Futures Trading Commission
Three Lafayette Centre, 1155 21st Street NW
Washington, DC 20581

VIA RULE-COMMENTS@SEC.GOV
Vanessa A. Countryman, Secretary
Securities and Exchange Commission
100 F Street NE
Washington, DC 20549-1090

RE: **Joint Request for Comment on Swap and Security-Based Swap Data Reporting**
91 FR 37877 (June 24, 2026) · CFTC RIN 3038-AF70 · SEC RIN 3235-AN78 · Release No. 34-105734 · File No. S7-2026-22

Dear Mr. Kirkpatrick and Ms. Countryman:

I

Introduction

Ariadne Dataworks Ltd. ("Ariadne") is grateful for the opportunity to comment on the Commissions' Joint Request for Comment on Swap and Security-Based Swap Data Reporting (the "Request"). Ariadne is a boutique consultancy focused on regulatory technology; the observations in this letter draw on our implementation work with open market data standards and on an archive we maintain of the swap data that repositories publicly disseminate, and we would be glad to share any of the underlying material with Commission staff.

The Request asks, across twenty-seven questions, how the reporting frameworks might be clarified, simplified, harmonized, and made more reliable. Most of our suggestions share a single premise, which we offer respectfully for the Commissions' consideration: **that the reporting framework itself could be published as structured data.**

Rule text creates obligations, but data quality is largely decided by four artifacts that sit downstream of rule text: the field dictionary, the validation logic, the lifecycle event vocabulary, and the reference data. Today those artifacts live as prose, PDF tables, spreadsheets, and per-repository interpretation, and every reporting counterparty, SDR, SBSDR, and vendor re-implements them independently. In our view, the divergence between those re-implementations accounts for much of what the Commissions observe as inaccurate, incomplete, and inconsistent data.

II

Summary of suggestions

We respectfully offer six suggestions, developed in the lettered sections that follow.

R1 **A single joint technical specification, maintained as a machine-readable artifact.** The Commissions may wish to consider maintaining one CFTC-SEC field dictionary as a versioned schema rather than a PDF, with any SEC-specific differences expressed as explicit deltas from the common baseline. (Section A)

R2 **Validation logic published in executable form, with a public conformance corpus.** Reporting outcomes would become testable before submission, and repository validations could be aligned to a single published set. (Section B)

R3 **A normative lifecycle event model with worked examples.** Explicit unit-of-report and linkage rules (UTI cardinality, package identifier, prior UTI), illustrated end to end for packages, strips, allocations, compression, and clearing. (Section C)

R4 **Reference data governed as part of the framework.** Static product attributes could move from per-trade reports into reference data, with freely accessible machine-readable records for every mandated identifier, completion of the commodity UPI designation, and alignment with the FDTA joint standards. (Section D)

R5 **A common public dissemination dialect with uniform access terms.** A normative dissemination schema, uniform free-access terms, a longer retention floor, dissemination timestamps, structured cap flags, and a price-forming indicator, while dissemination infrastructure remains with the repositories. (Section E)

R6 **Versioned convergence rather than a single cutover.** Dictionary first, then validations and corpus, then event model and reference data, with each phase gated on published data-quality metrics and SEC codification complete before the 2019 Compliance Statement expires on November 5, 2029. (Section F)

Section A

Harmonization may bind most durably at the technical layer

RESPONDS TO QUESTIONS 1 AND 4

The Request's own account of the last six years contains, we believe, its most encouraging harmonization lesson. Since 2019, SBS market participants have "universally relied" on the 2019 Compliance Statement and reported SBS transactions broadly in line with CFTC requirements. The instrument that achieved this in practice was a technical specification: firms harmonized because both regimes could be served by one artifact, one field dictionary, one set of formats, one pipeline. We respectfully suggest that this model is worth making permanent and deliberate.

> **Question 4**
> "Given the scheduled expiration of the 2019 Compliance Statement in 2029, should the SEC consider amendments to its SBS Reporting Rules to more fully harmonize those rules with the CFTC's swap reporting rules? Please explain."

In our view, yes. We would respectfully support amendments that harmonize the SBS Reporting Rules with the CFTC's swap reporting rules, with a technical specification identical to the CFTC's as the default position. Where a statutory difference proves unavoidable, it may be worth expressing that difference as an explicit, machine-readable delta from the common baseline: a named exception with its own justification, rather than a second document. Two prose documents maintained by two agencies may drift despite the best intentions; one artifact with two profiles is far harder to drift, because the differences are enumerated inside it.

> **Question 4.a (excerpted)**
> "If so, what should be the scope of such amendments? Should the SEC's and CFTC's rules be identical? … Should the SEC adopt a technical specification for reporting that is identical to the CFTC's technical specification? … [W]hat measures should the Commissions undertake to ensure that reporting rules remain harmonized over time?"

On the measures that could keep the rules harmonized over time, practices from modern specification maintenance may transfer well: maintaining the joint specification as a single version-controlled artifact under joint change governance; giving it semantic versioning and a published release calendar; having every change ship together with updates to a public test corpus (Section B); and publishing machine-readable diffs between versions so that every firm can see exactly what changed, mechanically. Harmonization would then become a property of the publication process itself, rather than an aspiration audited after the fact.

> **Question 1**
> "Which reporting requirements or data elements would benefit from harmonization to minimize or eliminate unnecessary inconsistencies in compliance obligations? Please identify and explain."

The elements that seem to us to benefit most are the ones both frameworks already share conceptually: identifiers, dates, notionals, prices, and action and event types, which are internationally aligned under the CPMI-IOSCO critical data elements and maintained under Regulatory Oversight Committee governance. We would, however, respectfully encourage the Commissions to look beyond data elements alone. The layer that today is not harmonized in any comparable form, and that in our experience drives much of the reconciliation cost, is validation and lifecycle logic: which submissions are accepted, which are rejected, and how lifecycle events must be sequenced. Harmonizing the dictionary while leaving validation logic to per-repository interpretation would preserve much of the operational complexity the Request asks about.

Section B

Machine-readable reporting logic and a public conformance corpus

RESPONDS TO QUESTIONS 15 AND 19

> **Question 19**
>
> "Should the Commissions integrate machine-readable rule structures or standardized reporting logic, and if so, how? What processes should the Commissions employ to maintain, update, and interpret such rule structures or logic over time? What are the benefits and risks associated with such approaches? What would the costs be, and would those costs be more or less than the costs under the current reporting frameworks?"

Our answer is yes, and we believe this may be the highest-leverage change available within the Request's scope; it is also the suggestion on which most of this letter's other suggestions rest. A design the Commissions may wish to consider has three components, in ascending order of importance:

1. **The field dictionary as a schema.** The technical specification could be published in a normative machine-readable format (a schema), with the PDF as a rendering of it rather than the other way around. Today the authoritative statement of up to 128 data elements is a PDF and spreadsheet that every implementer transcribes by hand.
2. **Validation rules as executable logic.** The validation rules could be expressed against that schema in an executable, openly licensed form, so that a reporting counterparty is able to run the exact rules a repository will apply before submitting.
3. **A public conformance corpus.** A versioned library of named, worked examples: complete submissions with their expected outcomes, accept and reject, spanning products, lifecycle sequences, and edge cases. The corpus is what would make the other two components trustworthy, because it is how any implementation proves it computes the same answers.

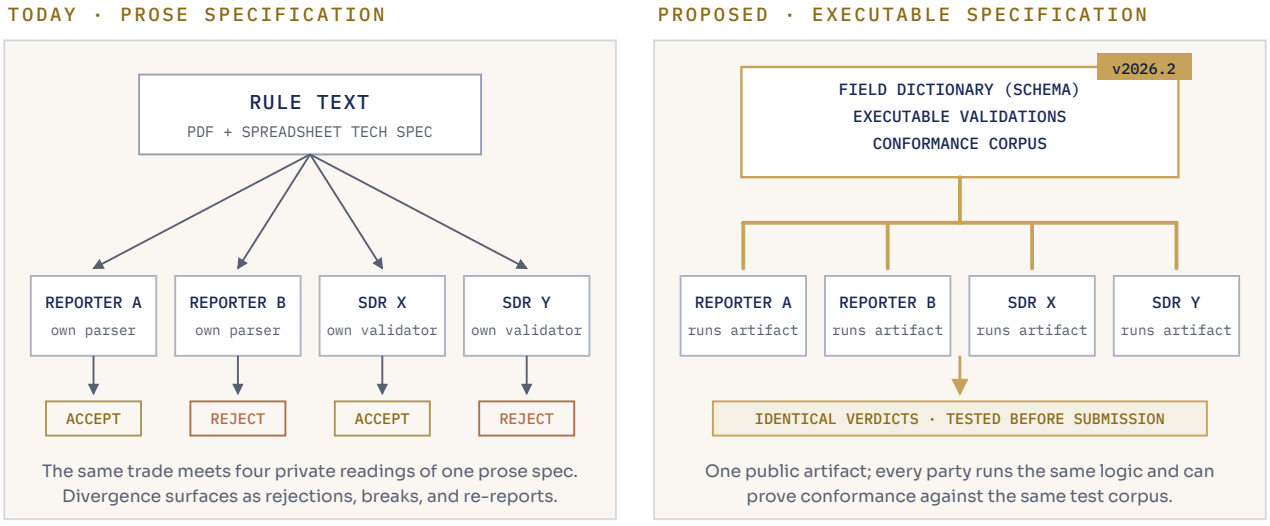

A shared artifact cannot drift apart; two separately maintained prose documents can.

FIGURE 1 · Divergence is structural when every party re-implements a prose specification privately, and is avoided by construction when every party runs the same published artifact.

None of this would be speculative. The post-rewrite regimes in the EU, UK, Japan, Australia, and Singapore disseminate ISO 20022 XML schemas whose message-level validation is deterministic by construction. On the industry side, ISDA's Digital Regulatory Reporting initiative already expresses the CFTC rewrite, and multiple further G20 regimes including the 2025 Canadian rewrite, as executable rules over the open-source Common Domain Model. The technology exists, is public, and is in production use. What seems to be missing is normativity: none of these artifacts is the authoritative statement of the rules, so every executable expression of them, however good, remains somebody's interpretation.

Our own implementation work with the CDM's executable validation rules suggests two lessons for the design Question 19 contemplates. First, the mechanics are unremarkable: hundreds of validation conditions compile and execute mechanically, and nothing about regulatory validation logic appears to require human interpretation at run time. Second, and more important, **the failure mode of machine-readable standards tends to be governance rather than technology.** Distributions drift, versions change frequently upstream, and two independent implementations of the same published standard can disagree silently for months; in our experience, nothing surfaces such divergence except executable tests. A machine-readable rule structure without a conformance corpus risks being a format rather than a standard.

On the processes for maintaining, updating, and interpreting such structures. We would respectfully suggest that the corpus can serve interpretation as well. When a participant asks how a particular fact pattern should be reported, the durable answer could be a new worked example added to the corpus: versioned, citable, and mechanically testable by every implementation from that release forward. Interpretive guidance would then accumulate as test cases rather than as prose letters, giving the Commissions a precise and comparatively low-cost instrument for correcting divergent practice as soon as it is observed.

On the risks the question raises. Three seem real to us. First, a defect in a normative artifact would propagate to everyone at once; the mitigations are the corpus itself, staged effective dates, and a fast published erratum channel, and the relevant comparison may be against today's baseline of many private defects propagating silently. Second, capture: the artifact would need to be public, free of charge, and openly licensed, rather than a product licensed by any vendor. Third, ossification: the change process would need to stay lightweight, with scheduled minor releases, or the artifact could lag the market it describes.

On costs. Today the same reporting logic is implemented privately by every reporting counterparty, every repository, and every vendor, and then reconciled after the fact through rejections, pairing breaks, and re-reports. A shared executable artifact would convert many implementations into many integrations of one implementation, and would move error discovery from after submission to before it. We would expect the net cost of this approach to be materially lower than under the current frameworks, with the savings concentrated where the Request's Section III points: rejection and correction workflows.

> **Question 15**
> "Should SDRs and SBSDRs take additional steps to validate reported transaction data to ensure accurate and high-quality public dissemination and regulatory reporting, and if so, for what specific data elements? Would such additional validations reduce the costs of addressing data quality issues? What specific additional validations should SDRs and SBSDRs add to ensure higher data quality?"

We would respectfully suggest that the more valuable step is alignment rather than addition: repository validations that match a published validation set exactly, identically across SDRs and SBSDRs, versioned in lockstep with the specification, and testable in advance against the corpus. The 2020 Final Rules narrowed SDR discretion over reported data elements after discretion had produced a lack of standardization; similar reasoning may apply to validation behavior. Additional validations seem likely to reduce industry costs only if they are uniform and pre-testable; divergent per-repository validations, however well intended, can themselves become a data quality problem, since a report accepted at one repository and rejected at another is not exhibiting a property of the report.

Section C

Lifecycle events and the unit of report

RESPONDS TO QUESTIONS 12, 16, AND 17

> **Question 16**
> "Which aspects of the current swap and SBS data reporting frameworks are most operationally complex or difficult to implement, particularly when considering the value of the information reported and the requirements of multiple reporting frameworks?"

In our experience, the most operationally complex aspect is not any individual data element. It is the mapping from one economic transaction to a set of reports over time: how a negotiated deal decomposes into reportable transactions, how many transaction identifiers it takes, how the pieces are linked, and which lifecycle message must follow which. The data elements are defined precisely; the unit of report often is not. That gap is invisible in the rule text and expensive on the ground, because it is resolved independently inside every firm.

> **Question 12 (excerpted)**
> "Are there data elements that, while conceptually useful, are persistently difficult to report consistently across registrants, asset classes, products, or market structures? If so, please identify and explain. What factors most commonly contribute to, and what transaction types (e.g., allocated trades) are most susceptible to, inaccuracies or inconsistencies in swap or SBS data reporting … ? Are these data elements static, and therefore could potentially be collected in a reference table, or do these elements change over time? … [A]re there practical steps the Commissions could take to improve data quality and transparency without significantly increasing reporting burdens?"

The transaction types that seem to us most susceptible to persistent inconsistency are the ones where the unit of report must be derived rather than read: package transactions, strips, allocated trades, and compression events. One finding from our research may illustrate the class. The word "strip" appears zero times in the CPMI-IOSCO UTI guidance, the CPMI-IOSCO CDE guidance, the

CFTC technical specification, and the ESMA reporting guidelines, based on corpus searches of each text; yet strips are an everyday unit of negotiation in commodities and rates. The correct treatment is derivable from first principles: the CPMI-IOSCO UTI guidance provides that "[w]here individual components of a package or strategy trade are reported separately, a different UTI should be used for each component," and the CFTC technical specification's package definition has each component assigned its own identifier and reported separately, linked by the package identifier with the negotiated price carried as the package transaction price. But every firm performs that derivation alone, and the two sides of one trade can derive differently.

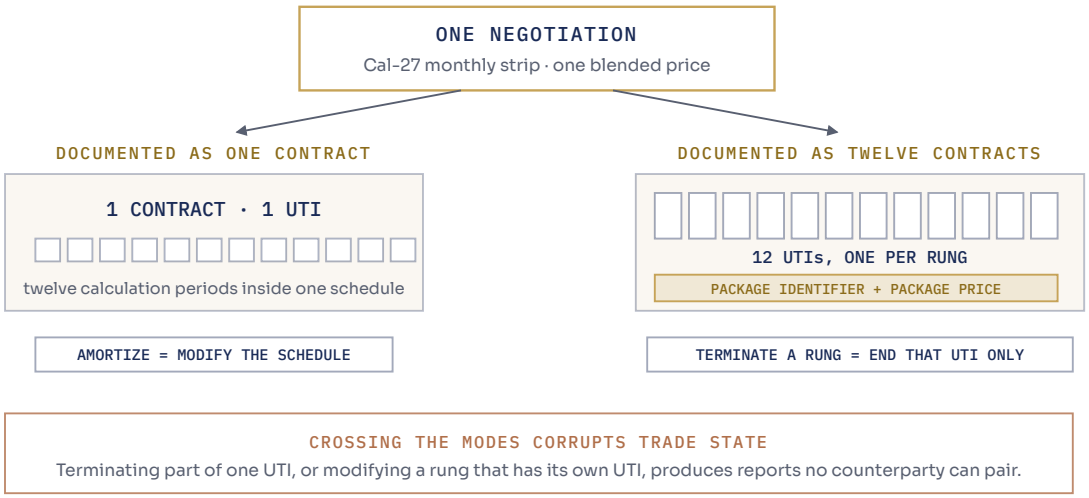

FIGURE 2 · One negotiated strip, two legitimate documentation forms, two different report structures. The framework does not state the rule that decides between them, so each firm derives it privately.

When the derivations disagree, the failure is structural rather than cosmetic: one side reports one contract and the other reports twelve, the reports can never pair, lifecycle events land on identifiers the other side does not know, and package indicators are misstated.

Precision at the data element level cannot, by itself, resolve a disagreement about how many reports there should be. Worked examples that state the unit of report for these transaction types, published as part of the corpus described in Section B, would in our view be among the most practical steps available under this question, and would improve data quality without adding reporting burden. And to the question's reference-table suggestion: yes, several persistently difficult elements are static product attributes, and as we later suggest in Section D, they may belong in reference data rather than in every report.

> **Question 17**
>
> "Are there particular validation rules or reporting obligations related to lifecycle events that could be simplified? Are there particular validation rules or reporting obligations related to lifecycle events where the current requirements are appropriate or inappropriate given the value of the information reported? Please explain."

We would respectfully caution against simplifying lifecycle validation by deletion: the complexity is real because the underlying lifecycles are real, and deleting validations tends to convert visible rejections into invisible data corruption. The simplification that seems most promising to us is to make the event model explicit: a small normative vocabulary of business events (new, modify, correct, terminate, error, and the clearing, compression, and allocation events) with defined state transitions, plus validations that assert the invariants mechanically: a package identifier present on all or none of a ticket's components; identifier uniqueness across them; package price consistency; and lifecycle events that stay in the mode the trade was born in, so that no one terminates "part of" one identifier or modifies a rung that has its own. Each of those invariants is machine-checkable at submission time; today they often go unchecked until two counterparties fail to reconcile, weeks later and at higher cost.

The event vocabulary might also usefully be anchored, with worked examples, to the execution trigger. "Execution" is defined as "an agreement by the parties, by any method, to the terms of a swap that legally binds the parties" (17 CFR 43.2, 45.1), and the CFTC's guidance distinguishes the affirmation that evidences execution from the later confirmation step. Real reporting chains run through an execution layer, a confirmation platform, and an internal booking system, and it is easy for reporting to drift toward whichever downstream event a firm's infrastructure finds convenient. Worked examples that name these stages, and fix the timestamps to the binding agreement, could quietly remove a systematic source of timing error.

Section D

Identifiers and reference data as governed parts of the framework

RESPONDS TO QUESTIONS 21 THROUGH 25

> **Question 21**
>
> "What limitations with respect to the information and supervisory or operational utility of the transaction data, if any, arise when the UPI or other standardized product identifiers or classification systems are used?"

We would draw attention to three limitations we have encountered. First, by design the UPI identifies a product template rather than an economic exposure: it carries no dates, no tenor, and no trade economics, so a two-year and a thirty-year swap on the same reference rate share one UPI. That is a sound design choice, but it means the UPI cannot bear the aggregation weight sometimes placed on it; transparency and supervisory analytics must join transaction-level fields regardless, and the disseminated data may be best designed on that assumption.

Second, coverage is incomplete: the commodity asset class still lacks a designated UPI in the United States, while Canadian rules will (tentatively) require commodity UPIs from January 2027;

completing the designation would remove an asymmetry in an otherwise standardized stack.

Third, access: programmatic access to UPI reference data is today mediated by registration, user agreements, and fee tiers, and for an identifier whose use the rules mandate, we would respectfully suggest that the reference record be retrievable freely, in bulk, in machine-readable form, by anyone.

> **Questions 22 and 24**
>
> "Are there instances where a data standard other than the UPI should be used to identify products? Please explain." · "Are there instances where a data standard other than the LEI should be used to identify counterparties? Please explain."

We do not believe an alternative standard should displace the UPI for products or the LEI for counterparties; the more useful direction, in our view, is completion and alignment rather than substitution. The Financial Data Transparency Act joint standards, finalized by nine agencies including both Commissions the day after this Request published (91 FR 38246, effective October 1, 2026), designate the LEI, the ISO 4914 UPI, and the ISO 10962 CFI, and deliberately adopt no common instrument identifier.

Swap and SBS reporting could build on exactly that stack and fill its gaps, rather than fork it. On the entity side, the LEI seems to us the strongest component of the current framework, and GLEIF's access model is much of the reason: free bulk downloads, open APIs, and versioned records. That model, and not merely the identifier, may be what deserves to be required of every mandated identifier system. The residual LEI difficulty, lapsed registrations, might best be handled deterministically, through a published validation policy stating how stale records are treated at report time, rather than through per-repository discretion.

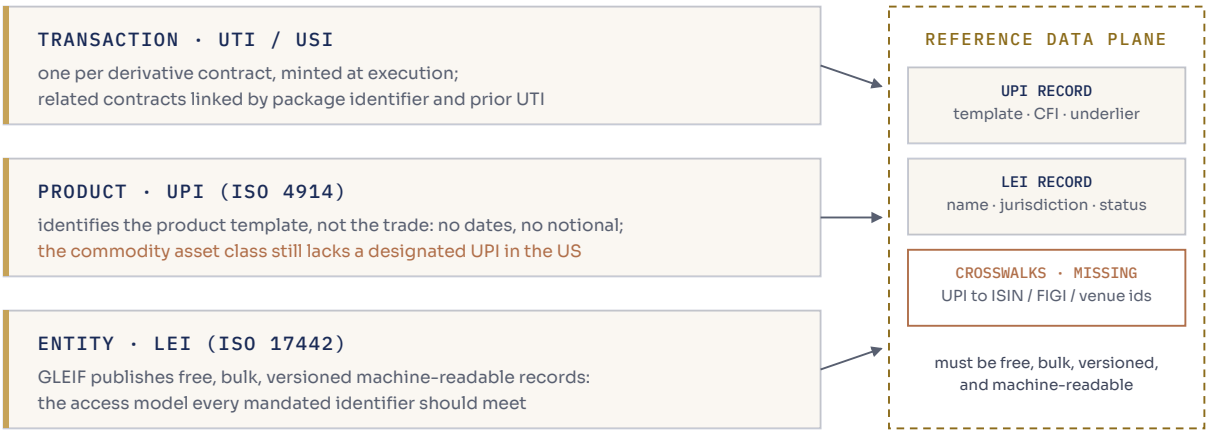

The FDTA joint standards (91 FR 38246) already designate LEI, UPI, and CFI across nine agencies and adopt no common instrument identifier. Swap reporting should build on that stack and fill its gaps, not fork it.

FIGURE 3 · The three identifier layers and the reference data plane they depend on. The plane would need to be free, bulk, versioned, and machine-readable; today the entity layer comes closest to that bar.

> **Question 23**
> "Are there additional opportunities to use standardized and static reference data elements to capture relevant attributes of swaps and SBS? Please explain."

The highest-value additional reference data may not be a new identifier but the connective tissue between the existing ones: authoritative crosswalks from UPI to classification (CFI), to underlier instrument identifiers (ISIN, FIGI), and to venue identifiers. In our experience, no authoritative, freely accessible crosswalk exists today; every data consumer rebuilds these mappings independently, and each rebuild differs slightly. The Commissions may wish to consider requiring, or themselves publishing, the crosswalks the reporting framework depends on.

> **Question 25 (excerpted)**
> "Is there information that is currently reported on a trade-by-trade basis that could more efficiently be captured through reference data? If so, please explain the benefits and drawbacks of using such reference data, including: (1) benefits or impacts to public transparency, usability, and accessibility of swap and SBS transaction and pricing data; (2) impediments to accessing reference data; (3) management of changes to the reference data that may occur over time; and (4) costs of maintaining and validating the reference data. …"

Information that is in fact a function of the product could, in our view, move to reference data. On the question's four factors: transparency seems likely to improve, since transaction records would become smaller and more consistent, and product attributes would resolve identically for every consumer instead of being re-keyed per trade.

Access impediments are, in our experience, the binding constraint and may deserve direct regulatory attention: reference data the framework depends on would need to be free, bulk, and machine-readable, with no contractual friction; today that bar is met at the entity layer, partially at the product layer, and not at all for crosswalks. Change management could follow standard reference data practice: effective-dated, versioned records, with reports resolving against the version in force at the event timestamp, and no silent mutation of history. And validation costs would likely fall, because validations would check a reference key once rather than re-validating repeated static attributes on millions of reports.

Section E

Public dissemination: observations from the public tape

RESPONDS TO QUESTIONS 13 AND 14 · EMPIRICAL OBSERVATIONS OFFERED PER SECTION III OF THE REQUEST

The observations below come from an archive we maintain of the publicly disseminated files across the CFTC, SEC, and Canadian regimes; we offer them as the data-driven input Section III of the Request invites, and can provide the underlying series to Commission staff.

- **Uniformity appears achievable, because one operator already achieves it.** The principal repository's cumulative public files share a single 110-column layout across all three regulatory regimes it serves. Cross-regime uniformity of dissemination format is not a hypothetical; it exists today within one SDR and stops at its edges.

- **Access practice does not always match the rule's text.** Part 43 requires publicly disseminated data to be "freely available and readily accessible to the public" in a "machine-readable electronic format," with instructions on how to download and search it (17 CFR 43.2, 43.3). In practice, one repository's public portal presents an explicit block page stating that automated access is prohibited by its website data terms of use, and another's public reports page is gated by IP address. For machine-readable data, automated retrieval is the point, and the Commissions may wish to clarify that website terms of use do not narrow the access the regulation grants.
- **Retention windows are short relative to their apparent cost.** The regulation obliges one year of availability; the principal repository commits to 366 days, and we observe roughly two years in practice. A full year of every cumulative file across all three regimes is on the order of 50 to 55 gigabytes compressed, so a retention floor of five years, or indefinite retention, appears operationally modest, and longer retention would preserve a public record that today disappears for anyone who has not archived it in time.
- **Some format details impose avoidable parsing burdens.** Cumulative files carry no dissemination timestamp (the publication date is inferable only from the filename); capped notionals are rendered as formatted strings such as "75,000,000+" rather than a value plus a structured cap flag (on a recent sample day roughly five percent of rates records were capped); and lifecycle churn dominates raw counts: a single day's cumulative equity file from one SBSDR contained 2,135,640 records, against tens of thousands for rates and credits across all three regimes combined, so identifying price-forming prints requires heuristics that every consumer invents independently.

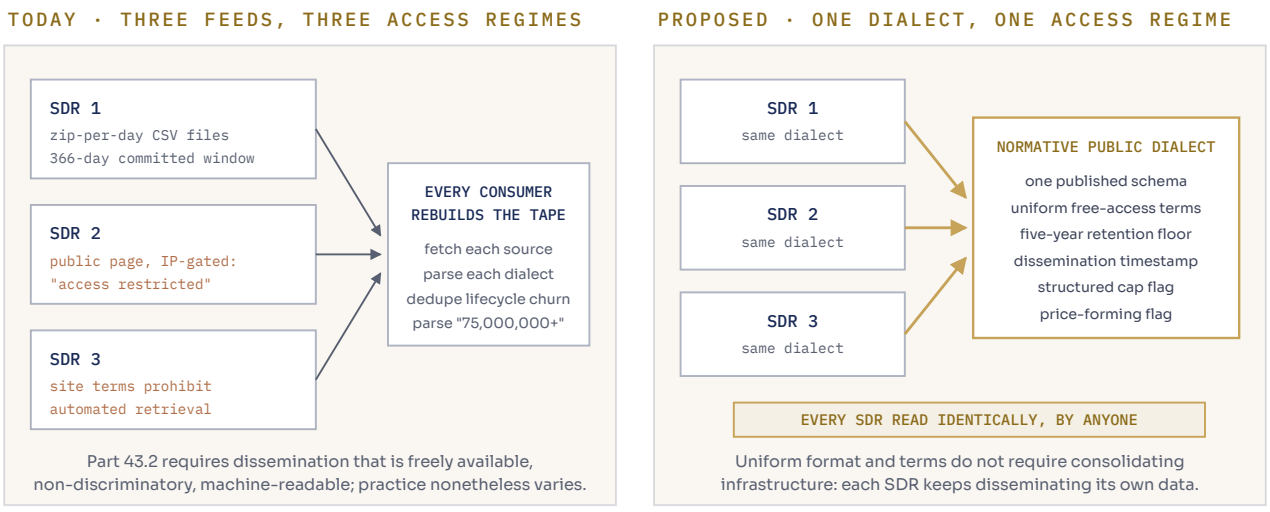

A full year of every cumulative file across the CFTC, SEC, and CSA regimes is roughly 50 GB compressed; longer retention appears inexpensive.

FIGURE 4 · The observed dissemination frictions are dialect, terms, and retention; none appears to require consolidating infrastructure to address.

> **Question 14**
>
> "Currently, each SDR and SBSDR, separately, publicly disseminates information about the transactions reported to it. Does this approach present challenges for price discovery or raise other concerns, or does it present advantages? Please explain."

Per-repository dissemination has genuine advantages: competition among repositories, resilience, and no single point of failure, and we would not suggest a consolidated tape operator. The frictions we observe are dialect, access terms, and retention, and all three seem addressable by specification: a normative dissemination dialect adopted as part of the joint technical specification (R1), uniform free-access terms consistent with Part 43's existing text, and a longer retention floor. Once every repository disseminates the same dialect under the same terms, anyone can consolidate the tape trivially, and the price discovery concerns the question raises would largely dissolve without new infrastructure. Should the Commissions wish to offer a consolidated public view, a thin index over uniform repository files might suffice.

> **Question 13**
> "Are there changes to reporting deadlines or methods of public dissemination that would improve the usefulness of the data?"

The reporting deadlines seem broadly right to us, and we do not suggest changing them. The gains, we believe, are in the method of dissemination: adding the dissemination timestamp to cumulative files, publishing capped values structurally, including a price-forming indicator derived by published logic rather than left to consumer heuristics, and versioning the dissemination schema alongside the reporting schema so the public tape evolves with the framework rather than lagging it.

Section F

Implementation: versioned convergence with published gates

RESPONDS TO QUESTIONS 26 AND 27

> **Question 26**
> "What factors should inform the appropriate timelines for any such changes, including the sequencing of changes affecting data standards, validation logic, and reference data frameworks of both the CFTC and SEC to minimize implementation risk?"

We would respectfully suggest sequencing by dependency. The field dictionary would come first, because everything else binds to it; executable validations and the conformance corpus second, because they are expressed against the dictionary; and the lifecycle event model, its worked examples, and the migration of static attributes to reference data third, because they consume both. Each phase could ship as a versioned artifact with published effective dates and adoption windows, with progression between phases gated on the metrics the Request itself lists in Section III: rejection and acceptance divergence across repositories, pairing rates, and correction frequency and latency. Gates of that kind would make implementation risk observable while it is happening rather than diagnosable afterward. The scheduled expiration of the 2019 Compliance Statement on November 5, 2029 may provide a natural outer deadline, since completing codification before then would avoid the market's current de facto harmonization lapsing into uncertainty.

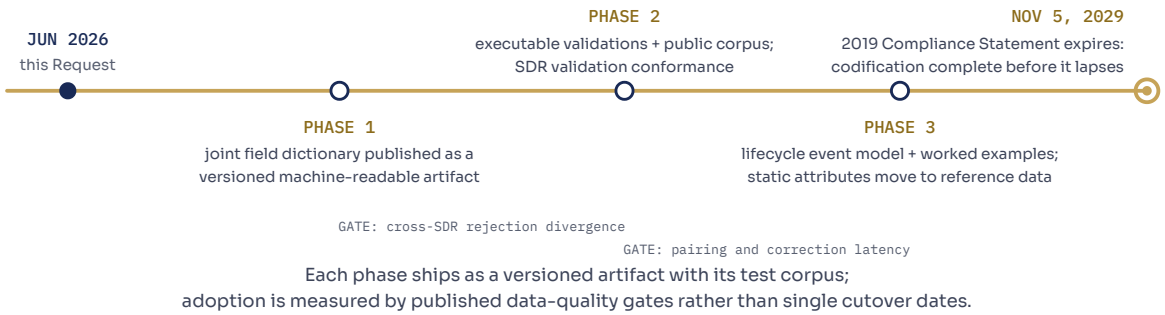

FIGURE 5 · Dependency-ordered phases, each a versioned artifact with its own corpus, gated by published data-quality metrics, completing before the 2019 Compliance Statement expires.

> Question 27
>
> "How can the Commissions structure implementation to minimize compliance costs and burdens?"

Four structural choices seem to us most likely to minimize cost. First, one shared executable implementation of reporting logic in place of hundreds of private ones; this may be the largest single saving available, and it would compound with every subsequent rule change. Second, versioned convergence rather than a single cutover: firms could adopt within windows aligned to their release cycles, against artifacts they can test months in advance. Third, reuse of what the market has already built and paid for: the CDE, the UPI, the LEI, ISO 20022 message standards, and the open-source CDM and DRR implementations, rather than new formats that restart the learning curve. Fourth, corpus-tested change: when every rule change ships with executable test cases, the per-change implementation risk that firms price into compliance budgets should fall, because conformance stops being a matter of interpretation.

III

Conclusion

The Request closes by observing that merely increasing the volume and granularity of data does not, by itself, enhance regulatory effectiveness. We agree, and would gently add that the same seems true of the volume of rule text. In our view, data integrity improves when every party to the reporting chain computes the same answer from the same inputs, and that becomes achievable when the dictionary, the validations, the event model, and the reference data are published as versioned, testable, machine-readable artifacts. None of the approaches suggested in this letter relies on speculative technology: the standards exist, the international precedents exist, and open-source implementations are in production. What appears to be missing is normativity, which only the Commissions can supply.

We are grateful for the breadth and openness of the questions the Request poses. We would welcome the opportunity to discuss any of the foregoing with Commission staff, and we are glad to share the underlying data series and analyses referenced in this letter.

Respectfully submitted,

Ariadne Dataworks Ltd.

usearіadne.com · contact@usearіadne.com

ANNEX A

Index of questions addressed

REQUEST QUESTIONS	SUBJECT	ADDRESSED IN
1, 4 (incl. 4.a)	Harmonization across frameworks; keeping rules harmonized over time	Section A; Suggestion R1
15, 19	Repository validations; machine-readable rule structures, their governance, risks, and costs	Section B; Suggestion R2
12, 16, 17	Persistently difficult reporting, operational complexity, lifecycle event validations	Section C; Suggestion R3
21, 22, 23, 24, 25	UPI and LEI limitations, alternative standards, reference data opportunities and factors	Section D; Suggestion R4
13, 14	Dissemination methods; per-repository public dissemination	Section E; Suggestion R5
26, 27	Implementation timelines, sequencing, and cost minimization	Section F; Suggestion R6

Questions not listed above fall outside our direct practice, and we respectfully defer to other commenters on them. Question text above is quoted verbatim from the Request, shortened only by omission (marked "…"), never by paraphrase. Empirical observations in Section E are offered in response to Section III of the Request (Request for Data); underlying series are available to Commission staff on request.