

A Frozen 2005 to 2025 EU ETS Data Product with Rule Based Validation and Cryptographic Integrity Commitments

Saim Atalay Keleştemur^{1,*}

¹ Department of International Trade and Finance, Istanbul Gedik University, Istanbul, Türkiye

ARTICLE INFO

Article history:

Received 7 July 2026

Received in revised form 21 August 2026

Accepted 4 October 2026

Available online 10 October 2026

Keywords:

Reproducible data product; EU ETS; Data integrity; Blockchain anchoring; Merkle verification

ABSTRACT

Compliance data from the European Union Emissions Trading System support enforcement, policy evaluation, and a growing body of research. The process of transforming registry records into published analytical series is not necessarily auditable. Scope harmonisation, historical corrections, and rolling releases make data reuse dependent on confidence in the processing pipeline. This study develops and evaluates a reproducible integrity workflow for a frozen country-by-activity-by-year-by-variable panel of verified emissions and allowance records covering two decades. The workflow combines canonical record identifiers, rule-based pre-acceptance screening, salted payload commitments, deterministic duplicate and replay protection, and Merkle root anchoring. It is assessed through tampering, duplication, and structural violation tests, a structured threat register, and repeated benchmarking on a local EVM chain. The results show that post-anchoring modifications, duplicate submissions, and malformed records are detected under the tested protocol. Merkle batching reduces on-chain commitment costs relative to individual anchoring while moving proof management off-chain. The study contributes a measured and reproducible integrity layer for EU ETS data. Cryptographic commitments make the anchored representation tamper-evident but do not establish the truth or completeness of the underlying emissions data. The system complements rather than replaces the Union Registry and the monitoring, reporting, and verification regime.

1. Introduction

The EU Emissions Trading System is the world's largest compliance market for greenhouse gas emissions. It includes more than 16000 stationary installations, 1600 aircraft operators and 2600 maritime operators whose reports are used for compliance enforcement and climate policy evaluation [1]. Carbon pricing revenues in national systems were over one hundred billion US dollars in 2025. Other jurisdictions are developing their own trading schemes [2]. The EU system has gone through four trading phases. Its permit auctions now transfer tens of billions of euros per year to the budgets of member states [3]. The system operates based on an annual cycle of monitoring, reporting and verification. Operators transmit verified emissions to competent authorities and surrender

* Corresponding author.

E-mail address: atalay@vitriol.ltd

<https://doi.org/10.65069/jessd21202624>

allowances in the Union Registry [4]. The legal framework provides rules for verification and the accreditation of verifiers [5]. The resulting records of verified emissions, allocation and surrender also contribute to a growing body of empirical research on carbon pricing [6]. Securities regulators analyze trading and auction data from the same market, which raises the bar for traceability of the underlying records [7]. This makes the reliability of these records a regulatory, analytical and scientific concern at the same time.

Public access to the records is mainly through the European Environment Agency data viewer aggregating registry transactions in country, activity and year series. The aggregation is useful but not neutral. Activity codes are harmonized, corrections of the scope applied and changes in coverage over time (aviation entered in 2012 and maritime transport with a phased obligation in 2024). Release notes and quality documentation are provided to the viewer [8]. However, the transformation from raw registry entries to the published series is not an auditable artifact in its own right. Reuse is not based on verifiable evidence for each step, but on confidence in the pipeline. This gap is crucial, especially in longitudinal studies. Only a stable vintage allows for a comparison of the market across two decades, as each refresh of the viewer can restate earlier years. None of the present distribution links a published analytical result to the exact registry snapshot from which it was derived. There is no guaranty that the filters and selection rules used by one study are the same as those used by another. A checksum ties a file to a moment in time but tells you nothing about the rules that created the file. Having a canonical definition of a record would remove this ambiguity, because every observation would have the same key no matter who extracted it. Abrell [9] demonstrates the value of closing this gap from a different angle. Without any ledger technology, he relationally reconstructs the transaction log and compiles documented data package of installations, accounts, holders, compliance records and transactions. This work is the nearest neighbor to this study and demonstrates the demand for registry level data products.

Distributed ledger technology has been talked about for the last decade for carbon markets. Already in the early literature, the application space was classified and the provenance and sustainability uses were surveyed in general terms [10, 11]. Soon the design space was constrained by the energy demand and throughput limits of public chains. The discussion on permissioned platforms shifted to who validates versus who reads [12, 13]. Mandaroux *et al.*, [14] describes an evaluation framework for a DLT based EU ETS in carbon markets including fraud reduction and audit efficiency. Schletz *et al.*, [15] develop decision framework and architecture for Paris Agreement carbon market mechanism, weighing transparency gains against governance and interoperability limits. Franke *et al.*, [16] translates the requirements of the same mechanism into platform selection criteria and compare a public with a permissioned chain as candidate designs. Zhao and Chan [17] question whether the cost of carbon trading of a blockchain is worth it. Their answer is a hybrid architecture, anchored on summaries, not every transaction, anticipating the batching question discussed in this study. The same pattern is reported in systematic reviews across the field. Provenance, transparency, and automation are highlighted in the literature, while oracle reliance, privacy, permissioning, legal validity, scalability, and cost are identified as open issues [18, 19]. Legal validity and governance are the least empirically studied areas because their examination requires an operating system rather than a proposal.

Also in the field are implemented systems. They show that construction is feasible but their domains are different than compliance registry records. In Hong Kong, a consortium chain based on Hyperledger Fabric has been piloted for the monitoring, reporting and verification of building photovoltaic carbon [20]. The ledger there is the immutable record behind the existing meters and it is site level measurement rather than aggregate registry publication. Real time monitoring and verification systems have been developed in industrial settings with their own credibility layers and

measured performance [21]. A simulated emission trading system on Hyperledger Fabric with encrypted firm data and learned auctions reports measured read and write rates [22]. Its object is settlement and auction privacy, not published registry records. Empirical work on tokenized voluntary credits is focused on price formation and liquidity, not on record integrity. It says that on-chain activity is still a thin layer over the registries of origin [23]. The chain itself does not avoid double counting in governance studies of bridging protocols but registry retirement and serial numbers do. They regard integrity as a regulated common resource, not as a property of the code [24]. The tokenization of voluntary credits also shows the promise and the challenge of bridging registries and ledgers [25]. But on one point these strands, which differ in approach and locus, agree. The registry remains the system of record. The ledger can at best make registry derived products more accountable. These streams together bring three open questions for compliance data. No reviewed study defines a canonical record identity that survives the harmonization boundaries of the viewer. None has an explicit correction and versioning boundary and combines duplicate and replay protection. No one reports the cost of implementing the commitment pattern on a historical panel of compliance scales. This workload is relevant and anchors annual batches of about one hundred and fifty aggregate records in a compliance cycle. No cited study was found that empirically demonstrates a complete record level workflow of the kind evaluated here on real EU ETS data.

This study fills this gap with a specific question set. The first question is whether rule based screening and cryptographic commitments can detect post anchoring alteration, malformed structure and duplicate or replayed submissions in real records. The second one asks whether a deterministic record identifier and Merkle proofs allow for later membership and immutability checks. The third is the difference between individual and Merkle anchoring in terms of throughput, latency, and on chain gas. The fourth asks what threats exist outside the system, especially source accuracy, completeness and monitoring methodology. The research analyzes the cost and the performance of the proposed commitment workflow under stated local chain conditions. The motivation is not to substitute the regulated EU ETS infrastructure. This is to make explicit the transformation and subsequent verification of a defined analytical data product, more reproducible and tamper evident.

The study contributes in three ways. The former is a known analytical data product. We build a panel of 3012 country–activity–year–variable observations (frozen in 2005–2025) from publicly available European Environment Agency data, sourced from the European Commission Union Registry. The selection of the scope, aggregation rules, definition of variables, the reconciliation procedure, checks for missingness and outlier screening for the construction of the panel are described in the manuscript. The second contribution is an integrity workflow that combines rule-based pre-acceptance screening, deterministic record identifiers, salted payload digests, duplicate and replay protection, and Merkle-root verification with an explicit versioning boundary for source corrections. The third contribution is an empirical evaluation, including tampering tests, a rule-by-rule structural validation, a threat register based on a published threat-modelling method [26], and repeated performance measurements of individual and Merkle-based anchoring on fresh local chains. These contributions collectively connect a documented EU ETS data processing procedure with an experimentally assessed cryptographic integrity workflow.

The system complements the Union Registry and the EU monitoring, reporting, and verification regime rather than replacing either. The remainder of the paper is organised as follows. Section 2 describes the data, the architecture, the threat model, and the benchmark protocol. Section 3 reports the results. Section 4 discusses their meaning for data quality and reuse.

2. Materials and Methods

2.1 Research Design

The study follows a design science research approach in the tradition of Hevner *et al.*, [27] and Peffers *et al.*, [28]. The artefact is a rule-based system that anchors verified emissions and allowance records on a distributed ledger and screens incoming records against integrity rules before acceptance. Three evaluation instruments accompany the artefact. A 21-year panel of registry data quantifies the economic exposure attached to data integrity. A threat register derived with the STRIDE method maps the ways in which the reporting pipeline can be subverted. A benchmark executed on a local Ethereum Virtual Machine chain measures the throughput, latency, and cost of anchoring under realistic load. The building blocks are established. Cryptographic hashing and Merkle trees are standard techniques, and the contribution lies in their integration with a regulated reporting pipeline and in the evaluation under realistic load.

2.2 Data Sources

The primary source is the European Commission Union Registry, distributed through the European Environment Agency data viewer, which assembles registry transactions, auction outcomes, and scope corrections. The July 2026 release of the viewer, downloaded on 27 September 2026, is used. The viewer reports verified emissions, freely allocated allowances, auctioned allowances, and surrendered allowances by country, main activity, and year. The coverage spans the years 2005 to 2025 and includes more than 16000 stationary installations, 1600 aircraft operators, and 2600 maritime operators [1]. The 2025 values combine reported figures with gap filled estimates where the verification cycle was still open at the time of download.

The secondary source is the World Bank Carbon Pricing Dashboard, which documents the number, type, price level, and coverage of carbon pricing instruments worldwide [2]. The dashboard supplies the global context against which the European system is positioned. Instrument level details are cross checked against the annual status report of the International Carbon Action Partnership [3].

2.3 Panel Construction and Aggregation Rules

The raw viewer file mixes detail rows with aggregate rows that carry the suffix 99 in the activity code, for example 20-99 for all stationary installations. Summing both levels double counts emissions. We therefore retain only the broadest aggregate per country, year, and variable. This is the 20-99 activity code for stationary installations, plus the 10-99 and 50-99 codes where aviation and maritime totals are reported. Rows from auction revenue accounts such as the Innovation Fund and the Modernisation Fund carry no emissions and are removed. The panel is stored in long form with one record per country, activity code, year, and variable. Five variables are kept, namely total allocated allowances, freely allocated allowances, auctioned allowances, verified emissions, and surrendered allowances. The resulting panel contains 3012 records, and the annual batch size rises from 114 records in 2005 to 158 records per year during 2013 to 2020. The panel was frozen on 27 September 2026 from the July 2026 vintage of the viewer, which a September 2026 refresh was noted to supersede. The frozen file with its checksum is provided in the supplementary material, so the results remain reproducible regardless of later refreshes.

2.4 Analytical Definitions

Verified emissions in year t are denoted E_t and total allowance supply is denoted A_t . The source reports A_t directly. It equals free plus auctioned allowances exactly from 2013 onward. For 2007 to 2012 a source correction category, labelled as corrections not reflected in the registry transactions, closes the identity, with a maximum gap of 39.7 Mt in 2007. The source correction category is

incorporated in the reported total allowance supply field A_t and is not stored as a separate blockchain record. The balance calculation uses the source reported total. All balance calculations in Figure 1 and Figure 2 and in Table 1 refer to stationary installations under the 20-99 activity code. Aviation and maritime transport enter only the 2024 composition of Figure 3. The annual allowance balance is

$$B_t = A_t - E_t \quad (1)$$

and the cumulative balance at the end of year t is

$$C_T = \sum_{t=2005}^T B_t \quad (2)$$

A positive balance denotes net new allowances entering circulation relative to verified emissions. The year-on-year change in verified emissions is

$$D_t = 100 \times E_t / E_{t-1} - 100 \quad (3)$$

The auction share of total allowance supply is

$$S_t = 100 \times Q_t / A_t \quad (4)$$

where Q_t is the auctioned quantity. Concentration is measured with the top five country share

$$CR5_t = \sum_{i=1}^5 s_{i,t} \quad (5)$$

where $s_{i,t}$ is the share of country i in total verified emissions in year t . Reporting anomalies are flagged with the modified z score

$$MZ = 0.6745 \times d / MAD \quad (6)$$

where d is the deviation of a country level change from the median of all pooled changes and MAD is the median absolute deviation of those changes. A threshold of 3.5 marks an outlier. Percentage changes enter symmetrically and are computed only where the preceding year is positive, so missing entries and new entrants are skipped. All quantities are expressed in gigatonnes of carbon dioxide equivalent unless stated otherwise. The series reports observed emissions under the scope of each year. Aviation enters in 2012, the United Kingdom leaves in 2021, and maritime transport enters with a phased obligation in 2024. The reported 52.2 % decline is therefore an observed decline in the reported series rather than a scope adjusted estimate.

2.5 Reference Architecture of the Traceability System

The artefact has five layers. The ingestion layer converts registry records into two digests. The record identifier is the Keccak 256 digest of the domain tag, the year, the country code, the activity code, and the variable code, with no value and no salt. The payload digest is the Keccak 256 digest of the same fields plus the numeric value and a random 128 bit salt. The salt is generated for every record and is stored with the payload outside the ledger. It prevents precomputation across records and forces a per record search. It does not provide confidentiality once the salt and the payload metadata are disclosed, since the value space of plausible emissions figures remains small. Payload confidentiality would require keyed hashing or encryption, which the artefact leaves to the deployment environment. Corrections in the source registry are handled outside the contract by anchoring a corrected release under a distinct release identifier, since the contract rejects repeat identifiers. Release versioning is a deployment choice and is not exercised in the benchmark.

The anchoring layer maps each record identifier to its payload digest in the registry contract. The contract reverts any attempt to anchor an existing identifier. Because the identifier excludes the value and the salt, a repeat submission of the same record with a fresh salt still collides on the

identifier and is rejected. This closes the double submission path that a per record salt would otherwise open. The verification layer recomputes both digests from the payload, queries the contract by identifier, and compares the returned digest with the recomputed payload digest.

Content rules execute off chain. The contract stores only digests and cannot inspect payload values, so the rule layer screens records in the ingestion gateway before hashing. It rejects values that exceed the physical capacity based bounds of the reporting activity and activity codes outside the regulated scope. The bounds derive from installation capacity and standard emission factors, so a legitimate surge in a small jurisdiction stays inside the permitted range.

The writer set is closed and small. It consists of the account of the registry operator and the accounts of the national registries. Individual installations do not write to the contract and keep their existing reporting workflow. This choice mirrors the structure of the published panel, in which national registries aggregate installation data before publication. The payload, the salts, and the Merkle proofs are held by the registry operator within the existing monitoring, reporting, and verification infrastructure. A copy can be mirrored on a content addressed store for redundancy.

The anchoring transaction records the submitting account and the inclusion timestamp as part of the ledger history, which supports non repudiation of submissions. Two anchoring modes are evaluated. In individual mode every digest is anchored as a separate transaction. In Merkle mode the digests are rolled into a binary Merkle tree [29] and only the root is anchored. The root for a set of n digests h_1 to h_n is obtained by pairwise hashing of adjacent digests until a single value remains. For an odd number of nodes at any tree level, the final unpaired digest is duplicated before pairwise hashing. The Merkle mode reduces on chain storage at the cost of an inclusion proof for each record, which is stored with the payload.

2.6 Threat Modelling Procedure

Threats were identified with the STRIDE method as systematised by Shostack [26]. Each threat was mapped to one stage of the compliance data life cycle, namely report preparation, verification, registry entry, publication, and retirement. Likelihood and impact were assigned by the authors with a predefined three level rubric. Likelihood represented the feasibility of the threat under the stated trust model, and impact represented the potential effect on data integrity or compliance accounting. The ratings are qualitative and are not empirical probability estimates. A mitigation was assigned from the architecture, and the residual risk after mitigation was recorded. The register in section 3.3 reports the outcome.

2.7 Anchoring Benchmark Design

The benchmark runs on a local single node Ethereum Virtual Machine chain with the Shanghai hard fork. The node was generated with Ganache version 7 at a fixed gas price of 20 Gwei. The registry contract was compiled with the Solidity compiler version 0.8.24 and deployed once. Three scenarios were executed. Scenario S1 submits the full 2005 to 2025 panel of 3012 real records in individual mode, of which 3004 are anchored after the guard rejects eight conflicting duplicates. Scenario S2 anchors 20000 randomly generated identifier and digest pairs, which serve only as payload for a synthetic transaction burst that stands in for a compliance deadline wave. Scenario S3 anchors the Merkle root of the real panel in a single transaction. Throughput is computed as

$$\theta = N/\Delta T \tag{7}$$

where N is the number of submitted transactions and ΔT is the wall clock time in seconds. The anchoring rate for accepted records is reported separately. Latency is measured per transaction as

the time from submission to confirmed receipt and reported at the 50th, 95th, and 99th percentiles. Cost is reported in gas units and in either at the fixed gas price, and converted to fiat currency as

$$C = G \times p \times P / 10^9 \quad (8)$$

where G is the gas used, p the gas price in Gwei, and P the ether price in US dollars. The metrics follow the data processing view of blockchain systems, in which throughput and latency are the primary dimensions of comparison [30].

2.8 Verification and Tamper Injection Protocol

After anchoring, both digests of every panel record were recomputed off chain. The contract was queried by record identifier, and the returned payload digest was compared with the recomputed value. A tampering experiment then altered the numeric value of 120 randomly drawn records by one tonne while keeping the identifier and the salt. A duplication experiment resubmitted an existing record identifier with a fresh salt and the same value, and the expected contract behaviour was a revert. A rule screening experiment applies five structural rules to every record before anchoring, covering schema, non negative values, regulated scope, admissible years, and admissible variable codes. The admissible year range extends to 2026 because the current open compliance year is admissible for reporting. Ten violations per rule were injected, fifty in total. A violation is counted as detected when the rule layer rejects the record before anchoring. The expected outcome was the rejection of all fifty injected records and of no genuine panel record. The detection rate is

$$R = 100 \times n_{det} / n_{inj} \quad (9)$$

where n_{inj} is the number of injected tampered records and n_{det} is the number that verification rejected.

3. Results

3.1 Verified Emissions and Allowance Balance

Verified emissions of stationary installations declined from 2014 Mt in 2005 to 963 Mt in 2025. The reduction over the full period is 52.2 %. Emissions peaked at 2165 Mt in 2007 and have fallen in almost every year since. The sharpest annual declines coincide with three episodes. The financial crisis of 2009 coincided with minus 11.3 %. The pandemic year 2020 coincided with minus 11.4 %, and the energy price shock of 2023 coincided with minus 16.5 %. The decline continued at minus 5.8 % in 2024 and minus 6.9 % in 2025. Figure 1 shows the path of verified emissions against the composition of allowance supply.

The annual balance between allocated allowances and verified emissions turned from surplus to deficit during the sample. The cumulative balance peaked at 837 Mt in 2013 and then fell steadily to minus 1408 Mt by the end of 2025. The decline in the cumulative balance coincided with the introduction of the market stability reserve in 2019 and is consistent with a reduction of the allowance surplus. The balance calculation itself does not quantify reserve absorption. The auction share of total allowance supply rose from 0.1 % in the first phase to 4 % in the second and 48 % in the third. It stands at 43 % in the fourth phase. Table 1 summarises the panel by trading phase, and Figure 2 traces the annual and cumulative balance.

adds 114 Mt of aviation and maritime emissions in 2025 that have no counterpart in 2005, and it fell by 46.5 %. The scope adjusted estimates that the viewer also publishes were not used. Table 2 summarises the data quality checks on the frozen panel.




Fig. 1. Verified emissions and allowance supply for stationary installations under activity code 20-99

Note: Bars show free allocation and auctioned allowances; the line shows verified emissions. Values are annual totals in Mt CO₂e. Source: European Commission Union Registry data viewer extract, July 2026 vintage, frozen on 27 September 2026; own calculations




Fig. 2. Annual and cumulative allowance balance for stationary installations

Note: Annual balance is source-reported total allowance supply minus verified emissions; cumulative balance is summed from 2005. Values are in Gt CO₂e. The balance does not quantify Market Stability Reserve absorption. Source: European Commission Union Registry data viewer extract, July 2026 vintage; own calculations.

Table 1

EU ETS panel summary by trading phase

Phase	Years	Verified, mean	Allowance supply, mean	Auction share, %	Cumulative balance, Mt
Phase 1	2005 to 2007	2072	2115	0.1	131
Phase 2	2008 to 2012	1942	2082	4	833
Phase 3	2013 to 2020	1700	1569	48	-211
Phase 4	2021 to 2025	1149	909	43	-1408

Note: Verified emissions and allowance supply are period means in Mt CO₂e. Auction share is total auctioned allowances divided by total allowance supply. Cumulative balance is C_T at the end of each phase. All quantities refer to stationary installations under activity code 20-99.

The reduction is sensitive to scope. Stationary installations fell by 52.2 %. The all sector series

Table 2

Data quality and reconciliation checks for the frozen EU ETS panel

Check	Result
Panel records after scope selection	3012
Conflicting duplicate identifier pairs	8
Missing verified emission values	0
Negative values	0
Out of scope activity codes	0
Years outside 2005 to 2025	0
Reconciliation of total against free plus auctioned allocation	Exact from 2013 to 2025, the 2007 to 2012 gap is closed by the source correction category, 39.7 Mt at most
2025 records carried as gap filled by the source	149
Robust outliers in country level changes	25 of 614

Note: The frozen panel covers 2005–2025 and retains the broadest aggregate record for each country, year, and variable. The source-reported total allowance supply includes the source correction category for 2007–2012; this category is not stored as a separate blockchain record. Gap-filled values are flagged as reported by the source. Robust outliers use the modified z-score threshold of 3.5.

3.2 Composition and Concentration of Reported Emissions

Combustion of fuels accounted for 589 Mt of verified emissions in 2024, close to half of the 1185 Mt total. Refining contributed 104 Mt, iron and steel 93 Mt, cement clinker 91 Mt, maritime transport 90 Mt, and aviation 62 Mt. Maritime transport entered the system in 2024 with a phased obligation. Chemicals, lime, ammonia, paper, glass, ceramics, ferrous metals, and hydrogen each contributed less than 35 Mt. Other activities contribute the remaining 35 Mt. Figure 3 shows the full composition.

Reported emissions remain concentrated in a small group of countries. Germany accounted for 273 Mt in 2024, Poland for 149 Mt, Italy for 101 Mt, Spain for 77 Mt, and France for 64 Mt. Together these five countries accounted for 664 Mt, equivalent to 64.2 % of stationary installation emissions and 56.0 % of the all sector total. The 2023 shock was unevenly distributed, as Figure 4 shows.

A robust outlier screen quantifies the reporting anomalies that an integrity system must surface. The 614 country-year changes in verified emissions have a median of minus 2.8 %. Twenty five changes, or 4.1 %, exceed the robust outlier bound. The largest movements concentrate in small jurisdictions, led by Liechtenstein with an increase of 241.7 % in 2017. The last observed United Kingdom value is in 2020, following the exit from the system in 2021.

Free allocation exceeded verified emissions in Luxembourg during 2024 and approached parity in Sweden, Iceland, and Finland. This pattern reflects the carbon leakage protection granted to exposed industries. The compliance reporting surface spans more than 16000 stationary installations, 1600 aircraft operators, and 2600 maritime operators, as Figure 5 summarises.




Fig. 3. Verified emissions by main activity in 2024

Note: Values are in Mt CO₂e. Activities follow the EEA main-activity categories; “Other activities” is the residual category. Aviation and maritime transport are included in this composition. Source: European Commission Union Registry data viewer extract, July 2026 vintage; own calculations.




Fig. 4. Country concentration and annual change in verified emissions

Note: The left panel shows the ten highest-emitting countries in 2024 and the top-five share of stationary-installation emissions. The right panel shows country-level changes from 2022 to 2023 for the countries displayed. Source: European Commission Union Registry data viewer extract, July 2026 vintage; own calculations.

Estonia fell by 37.4 %, Bulgaria by 36.2 %, and Portugal by 21.2 %, while Slovakia changed by only minus 2.4 %.




Fig. 5. Upstream EU ETS compliance reporting entities

Note: Counts describe the upstream monitoring, reporting, and verification workload and are not transaction counts for the evaluated ledger. The y-axis is logarithmic. Source: European Environment Agency EU ETS data viewer and associated metadata.

These entities define the upstream monitoring, reporting, and verification workload. The evaluated ledger does not receive one transaction per entity. It anchors the aggregated country by activity by year by variable records published by the registry.

3.3 Threat Register for the Reporting Pipeline

Table 3 presents the threat register. Each threat is classified with STRIDE, mapped to a life cycle stage, and rated for likelihood and impact before and after the assigned mitigation. The register singles out two integrity threats that dominate the risk profile. The first is the alteration of a verified figure between verification and publication. The second is the double submission of the same verified record, which inflates the apparent environmental outcome of the system. Both are detected under the tested protocol by the anchoring layer in combination with the duplicate guard of the contract.

Table 3

STRIDE Threat Register for the EU ETS Compliance Data Life Cycle

ID	STRIDE class	Life cycle stage	Threat	Likelihood	Impact
T1	Spoofing	Report preparation	A forged operator identity submits emissions data	L	H
T2	Tampering	Publication	A verified value is altered after verification	M	H
T3	Repudiation	Verification	An operator disputes a previously reported figure	M	M
T4	Information disclosure	Publication	Commercially sensitive installation data leak	M	M
T5	Denial of service	Registry entry	A deadline wave blocks anchoring	M	M
T6	Elevation of privilege	Verification	A malicious actor alters the rule layer	L	H

Table 3

Continued

ID	STRIDE class	Life cycle stage	Threat	Likelihood	Impact
T7	Double submission	Registry entry	The same verified record is submitted twice	M	H
T8	Oracle manipulation	Ingestion	A compromised feed injects false registry values	L	H
ID	Mitigation in the artefact		Residual		
T1	Closed writer set with signed submissions		L		
T2	Keccak 256 anchoring with public verification		L		
T3	Immutable digest anchored in a transaction that records the sender and the inclusion timestamp		L		
T4	The ledger stores the digest, the salt and payload remain off chain		M		
T5	Measured burst capacity with Merkle batching		L		
T6	Secure off chain execution environment with strict access control and authentication		M		
T7	Duplicate guard on the record identifier, repeat submissions revert		L		
T8	Out of scope of the on chain artefact, relies on off chain MRV audits		M		

Note: L, M, and H denote low, medium, and high. Likelihood and impact are qualitative ratings assigned under a predefined three-level rubric; residual risk is assessed under the stated architecture and trust model

3.4 Anchoring Throughput, Latency and Cost

The benchmark configuration is listed in Table 4. Scenario S1 submitted the complete 2005 to 2025 panel of 3012 records in individual mode and anchored 3004 of them. Scenario S2 anchored 20000 randomly generated identifier and digest pairs as a synthetic transaction burst. Scenario S3 anchored the Merkle root of the real panel in a single transaction, so its latency rests on one observation.

Table 4

Configuration of the Anchoring Benchmark

Component	Setting
Chain client	Ganache version 7 local node
Hard fork	Shanghai
Gas price	20 Gwei fixed
Contract compiler	Solidity 0.8.24
Registry contract	Deploy once, gas 250632
Record format	Record identifier plus salted payload digest
Mining mode	Auto mine on submission
Submission window	32 transactions in flight for S1, 64 for S2
Repeated runs	Five runs for S1 and S2 on fresh chains, S3 observed once
Host	Linux virtual machine, 2 cores, 4 GB memory
Scenario S1	3012 real records, individual mode
Scenario S2	20000 synthetic digests, individual mode
Scenario S3	Merkle root of the S1 panel, single transaction

Note: S1 uses the frozen panel in individual mode; S2 is a synthetic individual-mode burst; S3 anchors the S1 panel as one Merkle root. S1 and S2 are repeated on fresh chains, whereas S3 is observed once. TPS means transactions per second.

Table 5 reports the performance as the median across five repetitions on fresh chains. The median submission throughput of scenario S1 was 109.2 transactions per second with an interquartile range of 6.0. The anchoring throughput of accepted records was 108.9 records per second. Sequential probe transactions returned a median latency of 41 ms. The synthetic burst

sustained a median of 99.1 submitted transactions per second with an interquartile range of 0.6. Separately, sequential probe transactions returned a median latency of 42 ms and a median 95th percentile latency of 55 ms, and these probe values exclude queueing under concurrent submission. The observed throughput was limited by the signing and submission pipeline rather than by the chain. The single anchor consumed 52211 gas in every run, which confirms that gas is deterministic on the EVM. A representative single run preceded the validation experiments. It recorded submission throughputs of 96.8 and 91.4 transactions per second for S1 and S2, with p99 probe latencies of 61 and 80 ms. Scenario S3 was observed once, and its latency rests on a single observation. The source file contains eight zero valued duplicate rows for Czechia under total allocated allowances. The identifier guard rejected these during anchoring, and 3004 records were anchored. Figure 6 shows the probe latency distributions of the representative run.




Fig. 6. Probe latency distributions on the local benchmark chain

Note: Curves show empirical cumulative distributions for sequential S1 and S2 anchoring probes and verification calls. Anchoring latency is measured from submission to confirmed receipt; sequential probes exclude concurrent queueing. The curves represent the single benchmark run used for the displayed probe distributions

Table 5

Anchoring performance under the local benchmark configuration

Scenario	Submitted records	Anchored records	Median wall time	Median TPS	IQR of TPS	TPS range	Median probe p50 latency	Median probe p95 latency
S1 individual	3012	3004	27.6	109.2	6.0	99.6 to 112.5	41	53
S2 deadline burst	20000	20000	201.9	99.1	0.6	98.5 to 100.8	42	55
S3 Merkle root	1	1						

Note: S1 and S2 values are medians across five fresh-chain repetitions. Latency percentiles are median values from sequential probe transactions and therefore exclude queueing under concurrent submission. Submitted records include records rejected by the duplicate guard in S1. S3 is a single Merkle-root transaction.

The cost contrast between the two anchoring modes is large. Individual anchoring consumed 52211 gas per record, which sums to 156841844 gas for the anchored records, equivalent to 3.14 ether at the reference gas price. Merkle anchoring consumed 52211 gas in total, or 17.4 gas per anchored record, about 3004 times less than individual mode. Figure 7 displays the comparison.




Fig. 7. On-chain gas cost per anchored record for individual and Merkle anchoring

Note: Individual mode writes one transaction per record; Merkle mode writes one root and amortises its gas over the anchored records. The y-axis is logarithmic. The comparison covers on-chain gas only and excludes contract deployment and off-chain tree, proof-storage, and proof-verification costs.

The anchoring cost of the full panel in Merkle mode is therefore below the cost of two individual anchors. The comparison covers on chain gas only. Merkle mode additionally requires off chain tree construction, proof generation, proof storage, and proof verification, which the benchmark does not measure. Contract deployment is excluded from these per record figures.

3.5 Public network conditions and fiat cost

The local benchmark isolates contract cost and system throughput. Public networks add consensus delay and dynamic fees. We therefore measured the public endpoints on 27 September 2026. The median base fee on the Sepolia test network was 1.01 Gwei (Table 6). The median base plus priority fee on the Ethereum main network was 0.129 Gwei. The ether price was 2690 US dollars on 25 September 2026 [31]. Read calls over a public Sepolia endpoint returned in 618 ms at the median, against 108 ms on the local node.

Table 6

Fiat cost estimates for individual and Merkle anchoring

Fee scenario	Gas price, Gwei	Full panel, individual	Full panel, Merkle	One year, individual	One year, Merkle
Sepolia observed, 27 September 2026	1.01	426.12	0.14	21.42	0.14
Mainnet observed, 27 September 2026	0.129	54.43	0.02	2.74	0.02
Stress assumption	20	8438.09	2.81	424.15	2.81

Note: Costs are calculated in US dollars using the stated gas-price scenarios and the reference ether price. Merkle estimates assume one root transaction per batch. Contract deployment and off-chain tree construction, proof storage, and proof-verification costs are excluded.

The observed fee regime of 2026 makes anchoring inexpensive in either mode. The Merkle estimates assume one root transaction per annual batch, irrespective of batch size. The 2024 annual batch costs between 0.02 and 21.42 dollars depending on mode and network. The stress scenario separates the modes. At 20 Gwei, individual anchoring of the batch approaches 424.15 dollars while Merkle mode remains at 2.81 dollars. Gas consumption is deterministic on the EVM, so the cost driver is the fee regime rather than the contract. The write path on a public network follows the same protocol as the local benchmark, and a funded replication script is provided in the supplementary material.

3.6 Validation of the Implementation

Section 2.8 defines the correctness experiments for the implementation, including the rule screening protocol. These experiments are checks against the specification rather than empirical discoveries. The avalanche property of the hash function makes any single-value change detectable by construction. The duplicate guard is a contract invariant. Table 7 reports the outcomes of the six validation tests.

Table 7

Validation tests for the anchoring and rule-screening implementation

Experiment	Samples	Outcome
Tampered records rejected	120	120 rejected, detection rate 100 %
Duplicate identifier with fresh salt	1	Transaction reverted
Conflicting duplicate rows rejected	8	8 rejected at anchoring
Schema rule, injected violations	10	10 rejected, genuine rejections 0
Non negative rule, injected violations	10	10 rejected, genuine rejections 0
Scope rule, injected violations	10	10 rejected, genuine rejections 0
Year rule, injected violations	10	10 rejected, genuine rejections 0
Variable rule, injected violations	10	10 rejected, genuine rejections 0
Capacity based rules		Not evaluated
False record anchored and verified	1	Boundary confirmed

Note: Structural-rule tests use ten synthetic fixtures per rule. Capacity-based rules were specified but not evaluated because installation-level capacity data were outside the benchmark dataset. The false-record test is a boundary test and does not assess source accuracy.

The detection rates in Table 7 rest on ten synthetic fixtures per rule and do not establish real world false negative rates.

All 3004 anchored records verified against the contract after anchoring. Verification queries returned in 108 ms at the median and 194 ms at the 95th percentile on the local node. The double submission experiment confirms that the duplicate guard operates on the record identifier rather than on the payload digest. Resubmitting an existing identifier with a fresh salt reverted, so the guard holds even when the salt changes. The rule layer was exercised on the panel and on injected violations. It applies five structural rules that cover schema, non negative values, regulated scope, admissible years, and admissible variable codes. No panel record was rejected, and all 50 injected violations were rejected. The validation covers structural rules only. The capacity based plausibility rules were specified but not empirically evaluated, because installation level capacity data lie outside the benchmark dataset. The false record experiment marks the boundary of the guarantee. A well-formed record with a factually wrong value, one tonne for Germany for the 2026 compliance year, was anchored with its own identifier and salt and verified. The system therefore demonstrates detection of post anchoring modification. It does not prove accuracy, meaning that the source data was true. This outcome also settles the position of the artefact on false data injection, which threat

T8 records as out of scope for the on chain layer. Source accuracy remains the responsibility of the monitoring, reporting, and verification regime.

4. Discussion

This study presents the conversion of a frozen EU ETS panel into canonical records. Each record contains a country, an activity, a year and a variable and passes structural, cryptographic and reconciliation checks. These checks define the quality claim of the data product as reported in Table 2. Table 7 shows that the implementation was able to detect the injected post anchoring modifications and the structural rule violations under the tested protocol. The benchmark also demonstrates that the major efficiency improvement of Merkle anchoring is the alleviation of on-chain commitments while the proof management burden is still off-chain. This contrast is quantified in Tables 5 and 6. Overall, the findings suggest that the prepared EU ETS data product can be made tamper evident and auditable under defined rules.

The contribution is one step further than proposing architectures for carbon markets. Distributed ledger architectures for the EU ETS and for the mechanism of the Paris Agreement have been proposed already. Evaluation frameworks for these are available [14, 15, 17]. Field reviews confirm that most of the previous work is still conceptual, architectural or evaluative [19]. Here we bridge the gap identified in the introduction by an empirical treatment. It combines EU ETS specific record semantics with rule-based acceptance, deterministic duplicate protection, cryptographic commitments, Merkle verification, a threat register and measured execution overhead.

But the findings also say something about the data. Although the Union Registry behind the published aggregate series is a mature operational system, the published aggregate series is not a fixed research database. The series is affected by scope changes, activity code harmonization and historical corrections, which is why the all sector decline of 46.5% differs from the stationary decline of 52.2%. There is no bad data in the reconciliation residuals. They show that the allowance supply field includes correction categories, which are not present in the component fields [8, 32]. Abrell [9] arrived at a complementary conclusion by relationally reconstructing the transaction log. A frozen release with a checksum guarantees published results against subsequent viewer refreshes, something the rolling viewer can't do on its own. The data product shall be suitable for trend analysis, allowance and emission reconciliation, integrity and provenance studies, reproducible benchmarking and audit trail reconstruction.

The benchmark results must be read carefully. The median throughputs were stable over five repetitions and the interquartile ranges were narrow for both scenarios, as shown in Table 5. The gas per anchor was constant, confirming the deterministic part of the cost. The difference between the two anchoring modes is not accidental but structural. Individual anchoring Each record is written to the ledger. This has the highest on-chain cost. Merkle anchoring writes a single root and pushes tree construction, proof storage and proof verification off chain. The economic reading is based on the cost logic of the wider literature. The benefit of a ledger is relative to the value of the shared record relative to the operating cost [13, 17]. The benchmark describes the engineering trade off commitment pattern under controlled conditions. It confirms the feasibility of the pattern at the tested panel scale.

Two guaranty levels ought to be kept distinct. This distinction is the most important interpretative point of the study. Under the test protocol the system rejects post anchoring value alteration. It drops duplicate submissions of the same record identifier and retries with a new salt. It also rejects malformed records, negative values, activity codes out of scope, inadmissible years, inadmissible variable codes and records outside an anchored Merkle batch. These are the properties of the commitment layer as verified in Table 7. The system does not detect a source registry that returns an

incorrect value. It does not detect erroneous monitoring methodology, verifier error, or fictitious but well formed data. It does not identify missing records, source extraction that is incomplete, or country mapping that is incorrect. It also does not detect a malicious change to the rule layer or a compromised upstream feed, which the threat registers notes as residual risks. The boundary is directly demonstrated by the false record experiment of Section 3.6. The well-formed but factually incorrect record was anchored and verified. The experiment is the limit of the guaranty in concrete terms. Commitment is consistency with a represented anchoring. It does not demonstrate that the representation was factually correct pre-anchoring.

This study has several limitations. Firstly, the panel is based on selected aggregated records from the EEA data viewer and does not represent the full raw monitoring data at installation level. Second, changes to the scope of regulation affect historical comparability such as the inclusion of aviation and maritime transport and the United Kingdom's exit from the system. Third, values identified by the source as gap-filled in the 2025 release should not be interpreted as being epistemically equivalent to directly observed values. Fourth, the performance benchmark was executed on a local EVM node and does not set the performance of a production network. Fifth, the rule-validation experiments tested ten synthetic fixtures per rule and did not predict real world false-negative rates. Sixth, we specified capacity-based plausibility rules but did not empirically evaluate them due to a lack of installation-level capacity data in the benchmark dataset. Seventh, the correction and versioning from release to release was architecturally described but not exercised through a live correction cycle. Eighth, the Merkle comparison was measured on-chain gas only. The costs of off-chain tree construction, proof storage and proof verification were not quantified. Finally, reproducible reuse depends not only on the ledger commitment but also on the availability of the frozen panel, data dictionary, provenance manifest, scripts, test vectors, licence and persistent repository identifier, in line with the FAIR principles [33].

These limitations define a clear agenda for future work. Release-to-release correction tracking. Needs to be implemented and tested against successive vintages of EEA data with explicit supersession links and revision differences. Second, the capacity-based plausibility rules need to be tested against installation-level capacity data for both valid and invalid records. Third, the proposed pattern of commitment needs to be compared with signed manifests, cryptographic checksums and repository-based version control in terms of integrity coverage, storage requirements, verification effort and total operational cost. Fourth, the benchmark should be repeated in a production-like public or permissioned network, with realistic confirmation delays, fee variability, concurrent workloads, and failure-recovery scenarios. Finally, the release package should be deposited in a persistent repository with a DOI, an open license, machine-readable metadata and executable scripts so that independent users can rebuild the panel and verify the anchored commitments.

5. Conclusions

This study developed and evaluated a reproducible cryptographic provenance workflow for a frozen 2005–2025 EU ETS data product. The workflow integrates canonical country–activity–year–variable identifiers, salted payload commitments, duplicate and replay protection, Merkle anchoring, structural rule screening, and post-anchoring verification. Under the stated experimental protocol, it successfully identified the injected record modifications, duplicate submissions, and structural rule violations. Repeated benchmarking on fresh local chains quantified the throughput, latency, and on-chain cost trade-offs between individual and Merkle-based anchoring. Taken together, the results establish a practical integrity layer for transforming, committing, and auditing EU ETS records within the existing monitoring, reporting, verification, and Registry infrastructure. The principal contribution lies in the EU ETS-specific integration of deterministic record identity, rule-based acceptance,

cryptographic commitments, duplicate protection, Merkle verification, and empirical performance evaluation.

Data Availability

The EU ETS data used in this study are publicly available through the European Environment Agency's EU ETS Data Viewer and associated downloadable data products. These products provide country-, activity-, and year-level information on verified emissions, allowance allocation, auctioned allowances, and surrendered allowances derived from European Commission Union Registry records [1]. The analysis uses the July 2026 release accessed on 27 September 2026. The global carbon-pricing context is drawn from the World Bank's State and Trends of Carbon Pricing 2026 report [2], while complementary information on emissions-trading systems and their development is provided by the International Carbon Action Partnership's Emissions Trading Worldwide: Status Report 2026 [3]. The panel construction, aggregation rules, scope selection, variable definitions, reconciliation procedure, and analytical calculations are described in Sections 2.2–2.4.

Acknowledgement

This research was not funded by any grant.

Conflicts of Interest

The authors declare no conflicts of interest.

References

- [1] European Environment Agency. (2026). EU Emissions Trading System data viewer, July 2026 release. <https://www.eea.europa.eu/en/datahub>
- [2] World Bank. (2026). State and Trends of Carbon Pricing 2026. World Bank, Washington, DC. <https://doi.org/10.1596/978-1-4648-2348-0>
- [3] ICAP (2026). Emissions Trading Worldwide: Status Report 2026. Berlin: International Carbon Action Partnership. <https://icapcarbonaction.com/en/publications/emissions-trading-worldwide-icap-status-report-2026>
- [4] European Commission. Monitoring, reporting and verification of EU ETS emissions. (2026). https://climate.ec.europa.eu/areas-action/carbon-markets/eu-emissions-trading-system-eu-ets/monitoring-reporting-and-verification_en
- [5] European Commission. Commission Implementing Regulation (EU) 2018/2067 of 19 December 2018 on the verification of data and on the accreditation of verifiers. Official Journal of the European Union, L 334, 94-134. (2018). <https://eur-lex.europa.eu>
- [6] Dolphin, G., & Xiahou, Q. (2022). World carbon pricing database: sources and methods. Scientific Data, 9(1), 573. <https://doi.org/10.1038/s41597-022-01659-x>
- [7] European Securities and Markets Authority. EU Carbon Markets Report 2024. (2024). <https://www.esma.europa.eu>
- [8] European Environment Agency. ETC-CM EU-ETS data quality documentation. (2026). <https://www.eea.europa.eu/en/datahub>
- [9] [dataset] Abrell, J. (2026). EU Emissions Trading System (EU ETS) Data Package. Zenodo. <https://doi.org/10.5281/zenodo.20509231>
- [10] Casino, F., Dasaklis, T. K., & Patsakis, C. (2019). A systematic literature review of blockchain-based applications: Current status, classification and open issues. Telematics and informatics, 36, 55-81. <https://doi.org/10.1016/j.tele.2018.11.006>
- [11] Saberi, S., Kouhizadeh, M., Sarkis, J., & Shen, L. (2019). Blockchain technology and its relationships to sustainable supply chain management. International journal of production research, 57(7), 2117-2135. <https://doi.org/10.1080/00207543.2018.1533261>
- [12] Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., ... & Yellick, J. (2018). Hyperledger Fabric: a distributed operating system for permissioned blockchains. In Proceedings of the Thirteenth EuroSys Conference. Article No.: 30, 1 - 15 <https://doi.org/10.1145/3190508.3190538>
- [13] Sedlmeir, J., Buhl, H. U., Fridgen, G., & Keller, R. (2020). The Energy Consumption of Blockchain Technology: Beyond Myth. Business & Information Systems Engineering, 62(6), 599-608. <https://doi.org/10.1007/s12599-020-00656-x>

- [14] Mandaroux, R., Dong, C., & Li, G. (2021). A European emissions trading system powered by distributed ledger technology: An evaluation framework. *Sustainability*, 13(4), 2106. <https://doi.org/10.3390/su13042106>
- [15] Schletz, M., Franke, L. A., & Salomo, S. (2020). Blockchain application for the Paris agreement carbon market mechanism—a decision framework and architecture. *Sustainability*, 12(12), 5069. <https://doi.org/10.3390/su12125069>
- [16] Franke, L., Schletz, M., & Salomo, S. (2020). Designing a blockchain model for the Paris agreement's carbon market mechanism. *Sustainability*, 12(3), 1068. <https://doi.org/10.3390/su12031068>
- [17] Zhao, F., & Chan, W. K. (2020). When is blockchain worth it? A case study of carbon trading. *Energies*, 13(8), 1980. <https://doi.org/10.3390/en13081980>
- [18] Vilkov, A., & Tian, G. (2023). Blockchain's scope and purpose in carbon markets: A systematic literature review. *Sustainability*, 15(11), 8495. <https://doi.org/10.3390/su15118495>
- [19] Merlo, A. L., Mendonça, D. S., Santos, J., Carvalho, S. T., Guerra, R., & Brandão, D. (2025). Blockchain for the carbon market: a literature review. *Discover Environment*, 3(1), 68. <https://doi.org/10.1007/s44274-025-00260-4>
- [20] Luo, Y., Shen, J., Liang, H., Sun, L., & Dong, L. (2024). Carbon monitoring, reporting and verification (MRV) for cleaner built environment: Developing a solar photovoltaic blockchain tool and applications in Hong Kong's building sector. *Journal of Cleaner Production*, 471, 143456. <https://doi.org/10.1016/j.jclepro.2024.143456>
- [21] Li, Q., Shi, J., Li, W., Xiao, S., Song, K., Zhang, Y., ... & Lai, X. (2024). An efficient tool for real-time global carbon neutrality with credibility of delicacy management: A Modelx+ MRV+ O system. *Applied Energy*, 372, 123763. <https://doi.org/10.1016/j.apenergy.2024.123763>
- [22] Wang, Y. R., Ma, C., Ren, Y. S., & Narayan, S. (2023). Cooperative management of an emission trading system: a private governance and learned auction for a blockchain approach. *Financial Innovation*, 9(1), 122. <https://doi.org/10.1186/s40854-023-00547-6>
- [23] Tlili, H. (2025). Blockchain and tokenized carbon markets: Empirical evidence on market efficiency and transaction dynamics. *Sustainable Futures*, 10, 101109. <https://doi.org/10.1016/j.sftr.2025.101109>
- [24] Diaz-Valdivia, A., & Poblet-Balcell, M. (2025). The governance of the ReFi ecosystem: Integrity in voluntary carbon markets as a common resource. *International Journal of the Commons*, 19(1), 100-119. <https://doi.org/10.5334/ijc.1410>
- [25] Ballesteros-Rodríguez, A., De-Lucio, J., & Sicilia, M. Á. (2024). Tokenized carbon credits in voluntary carbon markets: the case of KlimaDAO. *Frontiers in Blockchain*, 7, 1474540. <https://doi.org/10.3389/fbloc.2024.1474540>
- [26] Shostack, A. (2014). *Threat Modeling: Designing for Security*. Wiley.
- [27] Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research1. *MIS quarterly*, 28(1), 75-106. <https://doi.org/10.2307/25148625>
- [28] Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of management information systems*, 24(3), 45-77. <https://doi.org/10.2753/MIS0742-1222240302>
- [29] Merkle, R. C. (1980). Protocols for public key cryptosystems. In *IEEE Symposium on Security and Privacy*, 122-134. <https://doi.org/10.1109/SP.1980.10006>
- [30] Dinh, T. T. A., Liu, R., Zhang, M., Chen, G., Ooi, B. C., & Wang, J. (2018). Untangling blockchain: A data processing view of blockchain systems. *IEEE transactions on knowledge and data engineering*, 30(7), 1366-1385. <https://doi.org/10.1109/TKDE.2017.2781227>
- [31] Fortune. Current price of Ethereum for September 25, 2026. (2026). <https://fortune.com>
- [32] European Commission. Union Registry. (2026). https://climate.ec.europa.eu/areas-action/carbon-markets/eu-emissions-trading-system-eu-ets/union-registry_en
- [33] Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., ... & Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data*, 3(1), 160018. <https://doi.org/10.1038/sdata.2016.18>