All You Need is Trust: A Longitudinal Analysis of Italy's OpenID Federation Journey

Tobias Hilbig
Munich University of Applied Sciences
tobias.hilbig@hm.edu

Erwin Kupris
Munich University of Applied Sciences
erwin.kupris@hm.edu

Thomas Schreck
Munich University of Applied Sciences
thomas.schreck@hm.edu

***Abstract*—As eIDAS 2.0 and the forthcoming European Digital Identity Wallet drive the evolution of digital identity, OpenID Federation (OIDFed) is emerging as a key technology for establishing trust in multilateral federations. Yet, little empirical evidence exists on how OIDFed performs at national scale and which operational pitfalls might emerge in practice. We conducted a longitudinal analysis of Italy's public administration federation, the first production-grade deployment based on OIDFed. Through daily scans between May 2024 and July 2025, we retrieved all publicly available metadata, examining the federation's topology, conformance, security posture, and evolution over time. Our analysis reveals an operational and mature deployment that experienced substantial growth, while also discovering topological irregularities, key material reuse, and conformance gaps. We disclosed all security issues to the respective stakeholders and publish our full dataset and tooling to support further research. Overall, our findings show that OIDFed can support a nation-scale production federation, while also identifying specification ambiguities and operational practices that must be addressed when deploying OIDFed in critical digital identity infrastructures.**

***Index Terms*—OpenID Federation, Federated Identity, EUDI Wallet, Critical Infrastructure, National Identity Systems**

1. Introduction

As Europe advances toward eIDAS 2.0 [42] and the rollout of the European Digital Identity (EUDI) Wallet [41], digital identity systems are undergoing significant transformations. A key component is the evolution of identity federations, i.e., trust frameworks that allow users to access multiple Relying Parties (RPs) using a single credential, while the RPs can verify identities through trusted Identity Providers (IdPs). Identity federations are widely adopted across critical infrastructure sectors such as government and healthcare, as well as in Research and Education (R&E).

The OpenID Federation (OIDFed) framework [27] is emerging as the foundation for such federations and as a modern alternative to the now-dated Security Assertion Markup Language (SAML) 2.0 standard [37]. OIDFed enables scalable and dynamic trust establishment between thousands of IdPs and RPs in OAuth 2.0 and OpenID Connect (OIDC)-based ecosystems. As a result, many SAML 2.0-based federations are expected to transition to OIDFed [35, 39, 20, 25]. During our measurement period, OIDFed was still evolving through Implementer's Drafts and working drafts. After our study period, in February 2026, it was standardized as OpenID Federation 1.0.

Many industry and public-sector stakeholders argue that the dynamic metadata model of OIDFed enables finer-grained, policy-aware trust decisions than conventional X.509 Public Key Infrastructures (PKIs) [9, 22, 36]. They propose to adopt OIDFed as an interoperable trust layer option in the forthcoming EUDI Wallet architecture. Members of the R&E community likewise argue that integrating the global inter-federation eduGAIN [21] with the EUDI Wallet via OIDFed would provide a decentralized, privacy-preserving trust backbone for millions of students and researchers [29]. Despite these proposals, OIDFed is not yet explicitly mentioned in the latest EUDI Wallet Architecture and Reference Framework (ARF) [8], although several technologies that are listed in the ARF declare compatibility with OIDFed [46, 34].

As an early adopter of OIDFed, Italy began the first nation-scale production deployment of OIDFed for its electronic identity system in late 2022 and has been expanding it ever since [20]. This system allows Italian citizens to authenticate digitally to thousands of public and private services using a single set of credentials, making it a piece of critical national infrastructure. In 2024, Italy also launched the IT-Wallet pilot, which uses OIDFed as its trust layer and already has over five million activated users and over eight million digitally verifiable documents [10].

These flagship deployments raise a practical question: how well does OIDFed perform in a live, nation-scale federation, and what lessons can be learned for future deployments? To answer this question, we conducted a longitudinal study of the Italian public administration federation by performing daily scans of all publicly available metadata over a 15-month period between May 2024 and July 2025. We assess the protocol's potential to securely manage trust in national-scale environments by evaluating the federation's usage of core features, policy mechanisms, and its conformance to the specification. We further analyze the evolution of the federation and classify security-relevant risks by interpreting conformance issues through a STRIDE-based analysis. We organize our findings around the following research questions:

- **RQ1** Which characteristics of the OIDFed specification are implemented by the Italian federation?
- **RQ2** How are advanced features such as metadata policies and constraints used by the federation?
- **RQ3** How well does the implementation conform to the employed specification draft version?
- **RQ4** How did the federation conformance and topology evolve between May 2024 and July 2025?

Our results show that the Italian federation has grown significantly, particularly through the integration of 7,000 entities related to the Ministry of Education and Merit. While our analysis confirms a mature and operational deployment, we identified conformance gaps and potential security issues. Our evaluation framework marked 12.28% of statements as invalid and 16.36% with warnings. We also observed topological irregularities that could potentially lead to policy violations during trust resolution, and notable reuse of key material. Security implications of these findings include spoofing, denial of service, and elevation of privileges. Finally, we disclosed all identified issues and provided detailed results to the relevant stakeholders, including the top-level authority of the federation.

This study makes three contributions. First, it provides the first longitudinal analysis of a nation-scale production deployment based on OIDFed, documenting how an evolving specification behaves in operational reality. Second, it identifies specification-level ambiguities and offers practical recommendations for organizations that intend to use OIDFed as a trust management solution for their digital identity infrastructures. Third, we publish reusable tooling and data to support future validation, monitoring, and comparative studies of OIDFed deployments [30].

This paper is organized as follows: Background information appears in Section 2, while Section 3 discusses related work. Our scanning methodology is presented in Section 4, followed by detailed results in Section 5. We discuss our findings, actionable recommendations, and directions for future work in Section 6. This study concludes with Section 7.

2. Background

To contextualize our research, we provide background information on federated identities and related standards. We also introduce OIDFed and the Italian CIE federation.

2.1. Federated Identity

As everyday life becomes increasingly digitized, users are required to manage an ever-growing number of accounts and credentials, leading to significant complexity and security challenges. Identity federations are an effective solution to reduce the need for multiple accounts by centralizing authentication. In its simplest form, known as a bilateral federation, an RP delegates identity management and verification to an IdP. When a user attempts to authenticate at such an RP, they are redirected to the IdP to verify their credentials. Upon successful authentication, the IdP issues a verifiable identity assertion and redirects the user back to the RP. Common examples include social logins, where a third-party IdP verifies user credentials and asserts their validity to the RP. In such cases, direct trust relationships and straightforward policy agreements between the involved parties are feasible.

For more complex ecosystems such as the R&E sector, multilateral federations spanning thousands of services have emerged. By relying on a unified trust framework mediated by third parties, these federations enable global collaboration. Users from one IdP can seamlessly authenticate at an RP operated by a different organization without the two entities needing a pre-existing trust relationship.

SAML 2.0 [37] has been the dominant standard for multilateral federations for nearly two decades. An example for the large-scale usage of SAML 2.0 is eduGAIN, the global R&E inter-federation that enables worldwide collaboration of academic institutions and their members [21]. By aggregating and publishing metadata files from national federations, eduGAIN acts as a trusted mediator. However, the now dated SAML 2.0 standard has limitations, including its static XML-based structure, lack of native support for modern web applications and APIs, and complex trust management.

While SAML 2.0 has been widely adopted across various sectors, many modern web applications rely on OIDC [43] and OAuth 2.0 for federated authentication and authorization. OIDC, an authentication layer built on top of OAuth 2.0, simplifies authentication by leveraging JSON Web Tokens (JWTs) [31] and RESTful APIs, making it well-suited for web, mobile, and cloud-based applications. In simpler federations, such as those found in enterprise environments, OIDC enables bilateral trust relationships between IdPs and RPs. Examples include corporate Single Sign-On (SSO) solutions, cloud-based IdPs, and social logins via providers like Apple, Google, and Microsoft. However, native OIDC lacks a standardized way for establishing dynamic trust in large, decentralized federations, leading to the development of OIDFed.

2.2. OpenID Federation

OIDFed [27] is an emerging protocol that addresses limitations in SAML 2.0 and traditional OIDC deployments by introducing a modern, scalable, and dynamic approach to enable multilateral federations. Initially designed as a trust framework for OIDC-based federations, the specification draft has since been generalized to support any kind of application protocol. OIDFed is expected to replace many SAML 2.0-based federations, particularly in the R&E sector [35, 39] and healthcare [25]. OIDFed 1.0 was published as an OpenID Final Specification in February 2026. During the period analyzed in this paper, however, the specification remained at the Implementer's Draft stage, with new working drafts released regularly. Appendix Table 2 summarizes the security-relevant fields that recur throughout our analysis.

Entities. OIDFed enables the dynamic discovery and verification of entities, e.g., IdPs and RPs, through a hierarchical trust architecture, as shown in the example in Figure 1. OIDFed's architecture is based on signed JWTs that can be composed to build cryptographically verifiable trust chains. Trust is established via a trusted third party, known as the Trust Anchor (TA), which serves as the root of a federation's hierarchical structure. Intermediate Entities (IEs) are subordinate to the TA, act as intermediaries, and issue verifiable statements about their subordinates. These subordinates can either be other IEs or Leaf Entities (LEs), which are the final link in the trust chain. Federated setups can support multiple TAs, and mutual trust between entities is possible as long as they share at least one common TA. Each entity provides its signing and possibly encryption keys as part of their metadata using JSON Web Key Sets (JWKSs). The TA's signing keys are provided to federation members via a secure, out-of-band way not described in the specification.

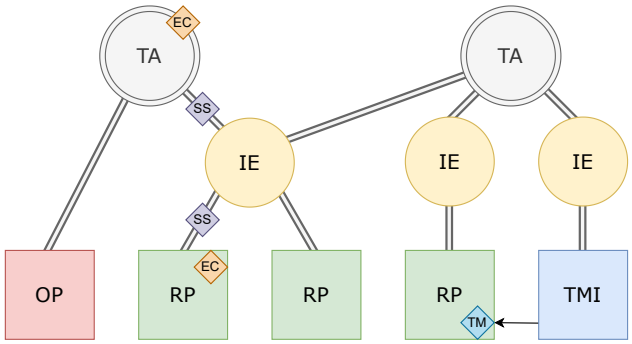

Figure 1. Example Architecture of an OIDFed-based Federation. For the first RP, a potential trust chain is shown. It includes ECs from the TA and the RP itself, as well as SSs about the IE and the RP.

Entity Statements. An Entity Statement (ES) is a signed JWT that provides information about an entity participating in a federation. Every ES contains both standard JWT claims [31], such as its issuer (`iss`), subject (`sub`), issue time (`iat`), and expiry time (`exp`), as well as OIDFed-specific claims depending on the kind of ES. An ES can either be a self-issued, self-signed Entity Configuration (EC), or a Subordinate Statement (SS) issued by a superior entity, such as an IE or TA, about one of its subordinates. In Figure 1, ECs are depicted as orange diamonds, while SSs are violet.

Trust Chains. A trust chain consists of a sequence of ESs, starting with the EC of an LE, ending with the EC of the TA, and including all intermediate SSs along the path. An example is shown in Figure 1, starting with the left TA and ending at the first RP. Since each entity includes their federation keys in its statements and references a superior entity that issues a verifiable SS about it, trust can be established dynamically. Each ES must be correctly signed and the referenced key material must match the signature for a trust chain to be considered valid.

Entity Types and Metadata. In OIDFed, the `metadata` claim in ESs contains entity-specific information essential for secure and interoperable federation operations. Every entity must provide Federation Entity (FE) metadata that is required to manage trust relationships. Depending on the role of the entity, other metadata is included in ESs. OpenID Providers (OPs) include information on authentication and token issuance. Similarly, an OpenID RP contains core OIDC parameters, informational metadata, and algorithmic requirements. In Figure 1, the entity types are colored as follows: green for RPs, yellow for IEs, red for OPs, and grey for TAs. Additionally, the specification defines OAuth 2.0-specific entity types, such as Authorization Server, Client, and Protected Resource.

Policies. By leveraging both technical and normative policies, OIDFed enables organizations to establish trustworthy federations that are both flexible and interoperable. The `metadata_policy` and `constraints` claims define rules for subordinate metadata and impose limitations on trust chains, enabling TAs and IEs to control the conditions under which their subordinates may participate in the federation. For example, they can restrict supported signing algorithms or specify allowed entity types. Both policies and constraints play a crucial role in maintaining consistency, security, and trust across the federation.

Trust Marks. Entities can provide Trust Marks (TMs) within their ECs to prove their conformance with specific trust, policy, and interoperability requirements. TMs have unique identifiers and serve as cryptographically verifiable statements that an entity meets well-defined sets of criteria. They are issued by trusted entities, often called Trust Mark Issuers (TMIs), as illustrated in Figure 1, in which the TMI, colored blue, issued a TM to the rightmost RP.

2.3. Italian Identity Federations

The Italian public administration operates three digital identity systems [1] that are converging on OIDFed for trust management. In the following, we briefly present systems relevant to OIDFed and their relation to our research.

The Carta di Identità Elettronica (Digital Identity Card, CIE), issued by the Ministero dell'Interno (Ministry of the Interior), serves as both a physical and digital identity document [16]. Its electronic component ensures secure online authentication for public services in Italy and the EU. CIE began its migration from SAML 2.0 to OIDC in late 2022 [20], and now actively relies on OIDFed for trust management. A technical project manager of Italy's Digital Transformation Department (DTD) confirmed that the CIE federation was based on OIDFed's draft version 22, remained aligned through draft 24, and awaited the finalized standard before updating further [13]. DTD material now indicates that CIE is planned to migrate to OIDFed 1.0 [12]. The CIE federation also defines its own technical specification [20] that outlines additional requirements for Italy's OIDFed implementations. For example, federation members are required to include the TMs issued by their respective authority in their EC. Entities are also required to expose an endpoint for trust chain resolution, and IEs are required to issue SSs for their subordinates.

The Sistema Pubblico di Identità Digitale (Public Digital Identity System, SPID), managed by the Agenzia per l'Italia Digitale (Agency for Digital Italy, AgID), allows citizens to access online services seamlessly using a username and password via a "Login with SPID" button [19]. It provides a unified digital identity, recognized across the EU, simplifying interactions with public administration and private entities. Although SPID also began migrating to OIDC in 2022 [20] and plans to use OIDFed for trust management [1], it only operates a single public OIDFed TA [45] which does not list any subordinates at the time of this study. Instead, SPID provides SAML 2.0 metadata for some residual entities through its API.

IT-Wallet, Italy's digital wallet solution, launched nationwide in 2024 inside the IO mobile app [11]. As of May 2025, it already serves over five million users and has issued over eight million verifiable digital credentials, including driver's license, health insurance cards, and European disability cards [10]. Even though the IT-Wallet specification explicitly references OIDFed [17], we could not find a publicly accessible TA.

While the three systems share the same national guidelines, they are operated independently and do not share a common trust hierarchy. Consequently, the longitudinal evaluation presented in this paper focuses on the CIE federation, which, at the time of our study, was the only deployment with a fully populated and publicly resolvable OIDFed topology suitable for comprehensive analysis.

3. Related Work

Research on OIDFed remains limited, presumably because the protocol was still in draft status during most early deployments and was only standardized recently. We therefore present both academic research and grey literature on OIDFed to position our study.

In 2020, Pöhn and Hommel [40] presented limitations and emerging solutions in the area of identity and access management, briefly discussing OIDFed as a promising alternative to SAML 2.0 in large-scale federations. They argued that OIDFed's dynamic, hierarchical metadata, published by each entity directly, is advantageous to SAML 2.0's reliance on aggregated metadata files published by federation operators. Kupris et al. [33] proposed a solution for securely automating the cumbersome IdP discovery process in multilateral federations using OIDFed. They utilized the trust management functionality of OIDFed and combined it with passkeys to prototype a more seamless and user-friendly alternative to standard "Where are you from?" procedures. While both publications highlight OIDFed's potential to advance federated identity, neither provides empirical analysis of real-world OIDFed deployments. Our longitudinal study aims to address this gap by evaluating the Italian OIDFed deployment and its conformity with the draft specification.

Complementing academic work, a growing body of grey literature explores the design, standardization, and early adoption of OIDFed. Connect2id provides an accessible technical overview of OIDFed's core concepts, positioning it as a solution to Internet-scale trust management [9]. Profiles for sector-specific deployments are also emerging. The OpenID Federation Wallet Architectures draft [15] applies OIDFed to digital wallet ecosystems by defining trust roles for Wallet Providers, Verifiers, and Credential Issuers. This work is evolving in parallel with the broader OpenID Digital Credentials protocols [46, 34]. In the education sector, the OpenID R&E Working Group [23] is profiling OIDC/OIDFed to modernize multilateral identity federations such as eduGAIN. In addition, Den Hertog et al. [29] argue that OIDFed could serve as a scalable and decentralized trust foundation for issuing academic credentials in the EUDI Wallet, leveraging the existing eduGAIN infrastructure.

In 2025, an OIDFed interoperability event [24] brought together 14 implementations from multiple countries to test draft 42 behavior and explore wallet integration scenarios. Italy's public-sector deployment has also contributed to the ecosystem's evolution. Federation operators presented early lessons from migrating CIE from SAML 2.0 to OIDFed [14]. Additionally, the open-source project OpenID Federation browser [18] allows exploring federations in real time. Released by Italy's DTD, it provides a web interface that discovers federation entities, resolves trust chains, and displays the result as an interactive graph.

In summary, the literature on OIDFed documents standardization progress, early tooling, and cross-sector pilots. However, it still lacks empirical, longitudinal measurements of real-world deployments. Our study addresses this gap by analyzing 15 months of data from Italy's national OIDFed federation, evaluating its topology, security posture, and conformance over time.

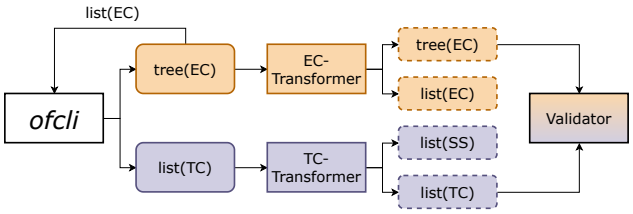

Figure 2. Sketch of our analysis pipeline. EC processing steps are orange, SS processing is violet. Decoded data is shown with dashed outlines.

4. Methodology

In this section, we describe the timeline of our research, our scanning and analysis pipeline, and the employed hardware and infrastructure. Additionally, we discuss the limitations of our study. Ethical considerations are provided at the end of the paper. We conducted daily scans between 2024-05-01 and 2025-07-31, covering 456 days. Scans ran daily at 03:30 UTC, from 2024-10-12 onward at 00:30 UTC. Each scan lasted between 21 minutes and seven hours, with a median of 2.25 hours. Scan duration varied due to changes in the federation size and performance optimizations of the employed tooling.

4.1. Procedure

We present a detailed description of our infrastructure for accessing and analyzing ESs. Figure 2 illustrates our custom-built processing pipeline, with EC-related steps colored orange and SS-related steps violet.

Retrieval. `ofcli` is an open-source tool capable of traversing federations based on OIDFed [26]. We slightly modified it to retrieve ECs and trust chains, relaxing certain checks to ensure that all data was fully captured for later analyses. The TA of the CIE federation [7] served as the entry point for our scanning. We retrieved all ECs of all entities directly or indirectly subordinate to the TA. Based on the set of entities from the most recent EC scan, we additionally retrieved all possible trust chains for each entity on 2025-07-31. We used a five minute timeout and did not consider validity and conformance at this step, allowing later, offline analysis. In addition, we implemented daily status emails to ensure timely notifications in case of issues, and contributed our changes to ofcli [26].

Transformation. We had to transform the raw output of ofcli in a format suitable for subsequent validation and analysis. Accordingly, we developed custom Python-based tooling capable of ingesting trees of ECs or sets of trust chains. It decodes all data, including nested JWTs like TMs, and validates all signatures using the PyJWT library. Any errors that occur during this process, e.g., signatures created with non-existent keys, are recorded for later validation. All decoded data is stored in a CSV file and a nested JSON to retain the federation's hierarchical structure.

Validation. We evaluate the conformity of the CIE federation following a systematic approach based on the requirements outlined in the specification draft. Our conformance tests are designed to be modular and easily adaptable, supporting both older and upcoming versions of the specification. As motivated in Section 2, we used draft version 24 as the baseline for our evaluation.

The validation framework classifies conformity issues by severity level, aligned with the specification's terminology. Non-compliance with absolute requirements and prohibitions resulted in errors, and violations related to recommendations in warnings, as defined in RFC 2119 [4].

Interpretation. We perform a retrospective threat modeling based on the STRIDE framework [28] to comprehensively interpret the security impact of observed conformance issues. STRIDE is widely used and offers a structured approach to threat categorization across six distinct threat types. It directly maps to the security properties most relevant in our setting, i.e., authenticity, integrity, availability, and privilege separation. Rather than modeling the entire ecosystem from first principles, we start from empirically observed misconfigurations and classify which STRIDE threats they may enable under realistic attacker models.

Hardware and Infrastructure. For our scanning and analysis tasks, we set up a virtual machine with eight CPU cores and 16 GiB of RAM. We employed Debian 12 on Kernel 6.1 with daily unattended security and software updates. The hypervisor was running on dual socket AMD EPYC 7702 64-Core CPUs. Access to the Internet with a symmetrical speed of 10 Gb/s was provided by the Leibniz Supercomputing Centre, and the DNS resolver by the Munich University of Applied Sciences.

4.2. Limitations

Our retrieval tooling produced broken output on 48 days (10.52%), resulting in small data gaps. Loops in the federation structure are discouraged and trust chain resolution must be aborted in such cases as of draft version 42. However, while such federation structures are problematic, they are technically possible. For this reason, some failures occurred during our scanning process due to entities incorrectly referencing themselves as subordinates. We consulted with one of the OIDFed draft authors and remediated this issue in `ofcli`. In addition, less than 1% of requests in each scan failed due to network-related errors, the majority of which were timeouts. All scanned entities complied with the draft's TLS requirement. Certificate validation relied on the Debian trust store, potentially excluding servers with invalid certificates. Our recorded set of entity data therefore represents a lower bound.

Until 2024-06-19, our tooling silently dropped JWTs with invalid `cty` header values due to strict validation in one of the libraries. We began recovering payload data of affected entities via debug logs on 2024-06-21, and from 2025-03-29 onward, we fully captured them. Our longitudinal conformity analysis differentiates between these sets to ensure temporal consistency and comparability.

Evaluating conformance against an evolving specification is inherently challenging. While the Italian deployment is based on draft version 24 [13], some components appear to have adopted features from more recent drafts. This ambiguity complicates analysis, as some of the issues flagged by our tooling may reflect version drift rather than real conformity issues. This highlights a broader limitation in assessing real-world OIDFed deployments while they are transitioning from drafts to the final specification: conformance results depend on draft version choice, which may not align with the implementation's intended state.

5. Results

We begin by presenting the federation's topology and evolution, followed by an analysis of claim prevalence and conformance to the draft. We then examine policy-related features and cryptographic aspects, before assessing the risks of the identified issues using the STRIDE framework.

The final scan on 2025-07-31 contains 12,789 ECs, 12,524 SSs, and 12,825 distinct TMs. Unless stated otherwise, all counts refer to this snapshot. For 326 entities (2.61%), no trust chains could be constructed. Therefore, some ECs have no matching SSs, while others have multiple. Table 4 in the Appendix summarizes our dataset.

5.1. Topology and Growth

For this topological analysis, we consider and compare our earliest and most recent scans, with their graphs depicted in Figure 3. The graphs are based on a top-down view of the federation, i.e., entities are connected to their superiors based on information provided by the superior's listing endpoint. The topology of an OIDFed compliant architecture directly results from the trust relations between the entities, with the draft advising implementers to build only tree structures. However, it is technically possible to construct directed acyclic graphs (DAGs) in which multiple paths exist between a TA and an LE. This is the case in the CIE federation, as some RPs are subordinate to an IE and the TA at the same time in both graphs.

We found various topological inconsistencies that violate the specification draft: Directly subordinate to the TA are 674 LEs that do not list the TA in their own `authority_hints`. In addition, 82 entities list the SPID TA [45] in their `authority_hints`, which only offers SAML 2.0 related metadata and does not participate in the CIE federation. Similarly, 35 entities list a non-existent pre-production server in their `authority_hints`. We also found 8,533 LEs that list the TA as their authority, while the TA does not consider them to be immediate subordinates. A single entity lists an IE as authoritative, while this IE does not consider the entity to be a subordinate. In 48 cases, ECs and SSs for the same entity contained the `entity_id` with and without a slash as a suffix. Finally, 261 LEs did not receive SSs from any of their superiors.

The topology also shows that the federation is a DAG with a height of two, i.e., at most one IE is located between the TA and any LE. This was the case during the whole scanning period. Most RPs are subordinate to an IE, while some are directly subordinate to the TA. The sole OP [6] is also a direct subordinate of the TA. All entities have a single entity type, except for the IE of the Ministry for Education and Merit [5], which acts as RP and FE.

We also determined the kinds of organizations that are represented within the CIE federation. These include public authorities, e.g., ministries, local and regional governments. Subordinate to these federal entities are various schools, other government bodies, public services, and communities with less than 2,000 residents. The geographic distribution of these entities shows nationwide coverage. In addition, we found many IT service providers, which make up more than half of IEs as of our latest scan. Nearly all of them offer digitalization services to the government, while some of them are even owned by public institutions.

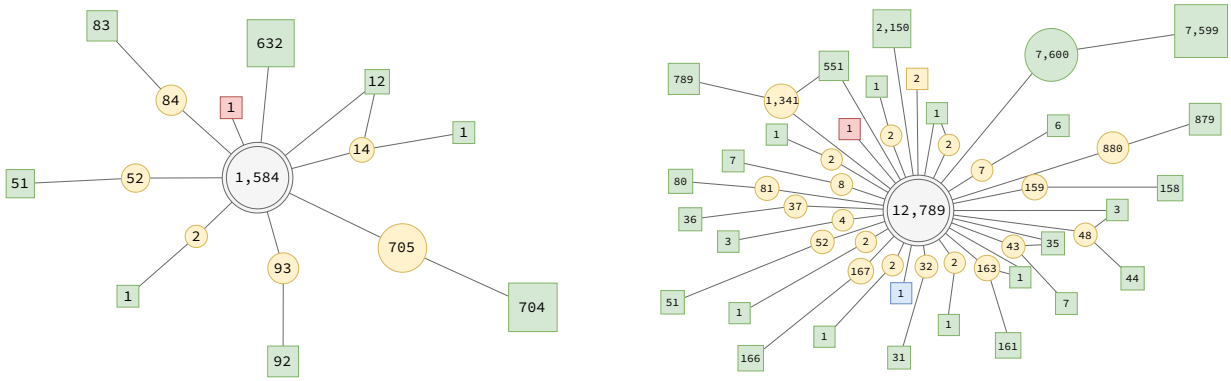

(a) First scan on 2024-05-01.

(b) Last scan on 2025-07-31.

Figure 3. The topology of the CIE federation depicted as a graph: (a) first and (b) last scan. Individual entities are depicted as circles; aggregated ones as squares. The TA is colored grey, IEs yellow, RPs green, OPs red, and TMIs blue. Each node is scaled and labeled with its number of sub-entities.

We consider all scans as a time series for the following longitudinal analysis. In Figure 4, we show the number of discovered entities by type over time. RPs are plotted on the primary Y-axis, while all other roles use the secondary axis. Entities with `cty` header claims, initially collected from debug logs starting 2024-06-20, are shown as dotted lines and fully included from 2025-03-29. In addition, detailed statistics on selected scans are shown in Table 3 in Appendix B. As shown in both graphics, a major addition to the federation occurred on 2025-03-08, where over 7,000 RPs subordinate to the IE of the Ministry of Education and Merit were added. Throughout our scanning period, the federation contained a single TA, one OP, and one TMI, with almost all leaves being RPs. The number of IEs increased from six to 21 between the first and last scan. The total number of entities increased substantially and steadily from 1,584 to 12,789. Finally, the topology of the federation did not change, with at most one IE between TA and LEs. However, structures with multiple trust paths appear more often in later scans.

Takeaway

CIE federation grew from 1,584 to 12,789 entities with nationwide coverage, minor topology issues, and inconsistent superior-subordinate relationships.

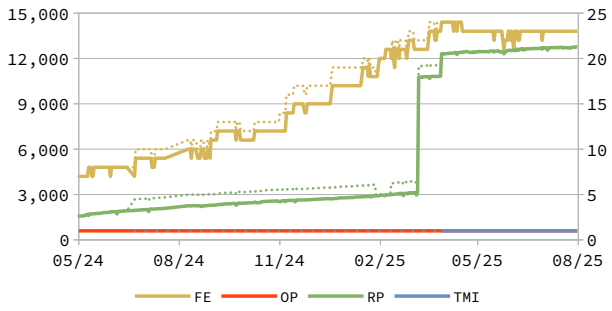

Figure 4. Entity counts over time by role. RPs use the left y-axis due to their amount; all others use the right y-axis. From 2024-06-20 onwards, entities with `cty` header claims collected from debug logs are depicted using dotted lines, and from 2025-03-29 onwards, they are fully included.

5.2. JWT Claims and Statement Types

This section analyzes ESs by statement type, beginning with JWT headers and core OIDFed claims. In addition, we examine claim prevalence and conformance of ESs.

JWT Headers. The analysis of JWT header claims across different statement types reveals consistent patterns. For ECs, SSs, and TMs, the mandatory header claims `typ`, `kid`, and `alg` are consistently applied across all instances. The non-standard `cty` claim, which we discuss in the limitations, is present in 1,538 ECs (12.03%) and 1,340 SSs (10.7%). Furthermore, TM headers show minor conformance issues with the `typ` claim: while most IEs use the correct value (`trust-mark+jwt`), a few apply non-standard alternatives, rendering these TMs invalid.

Core ES Claims and Lifetimes. All ESs include the required core claims `iss`, `sub`, `iat`, `exp`, and `jwks`, with no conformance issues apart from variations in their lifetimes. We compared `iat` and `exp` claims with the scan time to evaluate temporal validity. Most ECs (92.55%) were issued at the time of retrieval, and nearly all SSs (97.15%) showed similar behavior. A few outliers included timestamps far in the past or slightly in the future, potentially reflecting default values or clock drift. ES lifetimes are depicted in a histogram in Figure 5. The distribution of lifetimes aligns very well with the buckets, i.e., almost all values are at or near each bucket's upper end. While most ECs (60.79%) have a fixed lifetime of 20 years, lifetimes for the rest range from five minutes to over 1,900 years. The mean and median ES lifetimes were 12.37 and 19.64 years, respectively, and 95.07% of EC–SS pairs shared identical values. Most TMs (59.33%) are valid for 20 years, others range between a single day and a year. Only two omit the optional `exp` claim.

EC Claims. Core claims are consistently present in ECs: all include `metadata`, while `authority_hints` appears in all entities except the TA. As required by the CIE specification, 99.89% provide TMs. Over two-thirds of ECs are valid, while 20.94% raise warnings. 11.91% of ECs failed validation, with 88 of them containing two errors or more. Note that many invalid ECs also contain at least one warning. Errors primarily stem from embedded TMs, with additional issues including JWKS errors.

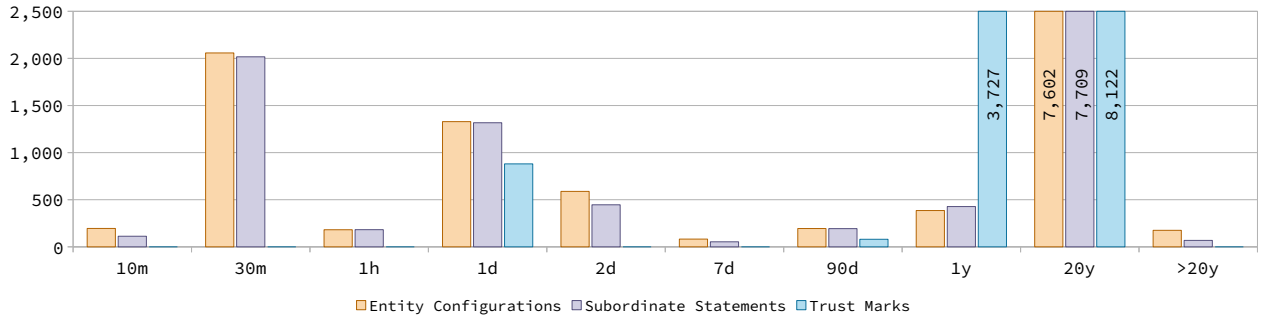

Figure 5. Chart depicting the instances of ES and TM lifetimes, aggregated by bucket. Note that the Y-axis of the 20y bucket is truncated.

We observed widespread use of non-thumbprint `kid` values, deviating from best practices [32]. The TA has the highest number of warnings among all ECs, with none of its three keys following this recommendation. Lastly, 44 ECs include unexpected or misspelled claims not documented in the federation's specification.

We visualize the number of ECs by conformance level over time in Figure 6. While the number of valid entities slowly rises, a notable share has warnings, with a smaller subset having errors as well. A significant rise in valid ECs occurs because of the successful onboarding of over 7,000 RPs. On 2025-03-29, improved tooling led to the inclusion of entities previously identified only through debug logs, which also contributed to a noticeable increase in ECs with errors. An additional, longitudinal analysis of errors and warnings is included in Figure 7 in Appendix E.

SS Claims. Beyond core claims, the use of optional fields in SSs varies widely. The `trust_marks` claim appears in all SSs, while the `metadata_policy` and `constraints` claims appear in 39.31% and 10.93% of SSs, respectively. A small number of SSs also include the `metadata` claim, defining OP and FE metadata objects. Only eight SSs include the optional `source_endpoint`. 75.65% of SSs are valid, 11.69% raise warnings, and 12.66% are invalid, with most errors again stemming from TM-related issues. We also observed repeated keys in 87 SSs and seven instances of invalid `constraints` values. About 10% of SSs issued by just two IEs and the TA do not follow the thumbprint-based kid convention. Other frequent warnings include unexpected or misplaced claims: over 1,000 SSs include the `fetch` claim outside the recommended location, and 77 incorrectly place `openid_relying_party` at the top level. In 31 cases, the non-standard metadata object `tokenSignedResponseAlgorithmId` appears instead, suggesting a typo or unsupported extension. Two SSs also contain a misspelled `metadata_policies` claim. These issues were isolated to three IEs, indicating localized misconfigurations.

Takeaway

ESs show consistent JWT headers but varied lifetimes, i.e., five minutes to 1,900 years. Two-thirds of ECs and three-quarters of SSs are valid. Trust Mark issues, non-thumbprint `kid` values, and unexpected claims appear in some instances.

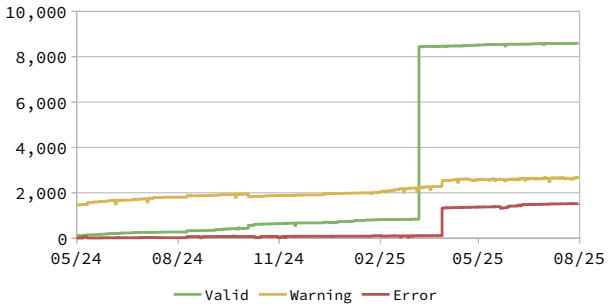

Figure 6. Chart depicting the number of ECs by validity over time. From 2025-03-29 onwards, entities with `cty` header claims are included.

5.3. Entity Type Prevalence and Conformance

In the following, we analyze prevalence and metadata conformance of entity types specific to OIDFed. Entity roles and types are both considered by our evaluation.

Federation Entities. All ECs in our latest scan correctly include FE metadata. Although the `resolve` endpoint is optional in the draft, it appears in all ECs, consistent with the CIE federation's stricter technical documentation [20]. Informational metadata is also widely present: `policy_uri`, `organization_name`, and `homepage_uri` appear in all ECs except the TA, while `contacts` and `logo_uri` are present in over 99% of cases each. This highlights a strong emphasis on organizational transparency and discoverability.

All 21 IEs and the TA also include the `list` and `fetch` endpoints, as required. Additionally, they expose a post-v24 TM-specific endpoint along with other features like `logo_uri`. This indicates that significant portions of the federation have adopted elements from newer OIDFed drafts, despite nominally implementing version 24.

Trust Anchor. The TA's EC is valid under draft version 24, though our conformance tests report five warnings. Three relate to its `kid` values not matching the JWK thumbprints, and another to the simultaneous presence of both `trust_marks_issuers` and its renamed variant `trust_mark_issuers`, which was introduced in draft 31. Again, this suggests that some components of the federation have evolved well beyond draft version 24. Additionally, the TA's EC defines the constraint `max_path_length = 1` to enforce a single intermediary per trust chain. While honored in practice, this claim is no longer permitted in ECs as of draft 35, and may require changes for future compatibility.

Intermediate Entities. In addition to the TA, 21 IEs list subordinates. While most expose only FE metadata, two also include RP metadata. During our scan, only 15 IEs actively issued SSs. Manual follow-up of the remaining IEs showed that two returned persistent server errors, while the other three appeared to be temporarily unavailable. One of these active IEs violates the draft by omitting trailing slashes from subordinate entity identifiers. Among the IE ECs, one is invalid due to TM-related errors, and four raise warnings: three for non-thumbprint kid values, and one for key reuse across roles. In addition, two entities include only `federation_entity` metadata but do not list subordinates, indicating possible inactivity.

When clustering conformance errors by superior IE, a clear pattern emerges: leaf entities subordinate to a single IE account for nearly all TM-related errors, which dominate the overall error set. Fewer than 200 EC and 250 SS errors originate from other parts of the federation. Only five IEs have subordinates with errors, while ten maintain fully conformant subtrees. While warnings are more widespread, they also cluster within specific IEs.

OpenID Provider. The federation's sole OP includes standard OIDC metadata and selected OIDFed-specific claims. It supports automatic client registration and disables manual registration. All major endpoints are present, and supported features include strong cryptographic algorithms, `private_key_jwt` authentication, as well as the `authorization_code` and `refresh_token` grant types. In addition to SPID assurance levels, attributes specific to the Italian eID, such as `fiscal_number` and `idANPR`, are also present. The OP's EC is valid with four warnings: its federation key is reused in the OP metadata, where another key is also present. Both keys lack a thumbprint-based `kid` value.

Trust Mark Issuer. The federation includes one TMI, correctly listed in the TA's `trust_mark_issuers` claim. It publishes a valid EC but appears inactive: its status endpoint returns HTTP 404, and it issues no TMs. Instead, the TA and IEs issue TMs directly. Although the TMI role was deprecated in draft 25, its metadata includes both the deprecated and updated claim variants, likely to establish forward compatibility, mirroring similar behavior observed in the TA. This reinforces our observation that the federation selectively adopted features from newer drafts.

Relying Parties. Nearly all ECs include consistent RP metadata, especially for core OIDC parameters like `redirect_uris`, `response_types`, `grant_types`, and `client_id`, which are present in all instances. Closely following in prevalence are ID token–related claims (88.43%), as well as the `application_type` claim (89.31%). Informational metadata such as `organization_name` (17.99%), `contacts` (14.13%), `subject_type` (9.21%), and `scope` (1.37%) appear less frequently. As expected given their volume, most EC errors originate from RPs.

Takeaway

Entities show strong metadata consistency; few IEs cause most validation issues. Some entities include claims from newer OIDFed versions, indicating a draft drift.

5.4. Policy Analyses

We now analyze policy-related features of the federation. This includes the `metadata_policy` and `constraints` claims, as well as the use of TMs.

Metadata Policy. Only 39.31% of SSs contain the `metadata_policy` claim, mostly (98.42%) defining policies for `openid_relying_party` metadata. These policies regulate key OIDC parameters such as supported grant types, response types, encryption and signing algorithms, and client authentication methods. They often specify exact or restricted value sets, and many mark parameters as `essential`, meaning that deviations would render the RP metadata invalid. A particularly important case is the SSs issued by the TA to the IEs. These include highly detailed policies directed at RPs, indicating centralized policy management. This suggests a hierarchical governance model within the federation. Supporting this, 97.85% of `metadata_policy` claims include `jwks`, suggesting that key material is centrally managed and propagated through the trust hierarchy.

The TA also issues one policy specifically targeting the federation's sole OP. This policy enforces strict requirements for nearly all OP metadata fields, including endpoints, algorithms, scopes, and client registration behavior, with many marked as essential. The OP's published EC metadata closely mirrors this policy, suggesting tight coordination between the TA and the OP.

Constraints. Most `constraints` claims in SSs (97.88%) contain an empty JSON object ({}), while seven use the object encoded as a literal string ("{}"). The frequent use of empty values is noteworthy, as it indicates that no effective constraints are being applied. All SSs with such ineffective constraints were issued by the same two IEs. Only the SSs issued by the TA about the IEs define a meaningful value for the claim by setting `allowed_leaf_entity_types` to `openid_relying_party`. As previously detailed, the TA also includes a `max_path_length` constraint. In addition, a single RP includes the `constraints` claim, but incorrectly has its metadata object embedded within. While no naming constraints exist, the domain name of an IE is employed by its subordinates in almost all cases.

Trust Marks. Although the current OIDFed draft discourages the use of TMs in SSs, the implemented version 24 permits this. Accordingly, we analyzed 12,825 TMs in ECs and 12,478 in SSs. Almost all TMs have valid signatures, with only 14 invalid cases in ECs. TM values are highly consistent across statement types: 12,444 entities have identical TMs in both ECs and SSs, with only a few TMs found exclusively in one or the other. Due to this high overlap, we focus our analysis on the 12,811 valid, EC-derived TMs. Nearly all entities include a single TM, as required by the CIE specification, while a small number include a second (69) or none (10). Most TMs (59.31%) were issued at the time of retrieval, while 9.27% were expired. The expired TMs are issued by the TA and four IEs, with three of them only issuing expired ones.

The federation uses different TM IDs in a structured way. The TA defines TM IDs for IEs, RPs, OPs, and OAuth resources. For each entity type, a public and private variant exists, allowing for eight IDs in total. Within the federation, five of those are in use; an additional TM ID related to a

pre-production server also exists. We found no instances of incorrect ID usage, e.g., RPs presenting TMs intended for IEs. Since almost all entities are RPs, the most common TM ID (99.66%) is `relying_party/public`.

A few isolated anomalies exist: The Ministry of Education and Merit functions as both IE and RP, but only offers an `intermediate/public` TM. All but eight entities listed in the TA's metadata are actively providing their subordinates with TMs. The TA itself issued TMs to only 2,200 of its 2,766 direct subordinates, while IEs in general issue TMs to all their subordinates. Finally, one IE omitted a TM for a single subordinate, and one entity embedded a TM in its EC that appeared to be an SS.

Takeaway

Federation shows centralized, hierarchical policy control, minimal constraints enforcement, and consistent yet partially incomplete Trust Mark usage.

5.5. Cryptography

In this section, we evaluate the cryptographic characteristics of the federation. We focus on the integrity of JWTs, the structure of JWKSs, and reuse of key material.

JWTs. All ESs and all but 14 TMs in our latest scan passed signature validation. Of the invalid TMs, seven were issued by a now-defunct pre-production server, making verification infeasible. Prior to 2025-04-09, the OP's TM was also invalid, as it was self-signed while claiming to be issued by the TA. Finally, all ECs and TMs were signed using the RS256 algorithm, as recommended by the specification draft.

JWKS. Nearly all entities in the federation publish consistent key sets in both their ECs and SSs, as required by the specification. The only exception adds one extra key in the SS, but reuses the same signing key across both statements. Therefore, our analysis focuses on 16,445 keys from ECs. Most entities (97.02%) publish at least one signing key, and 28.2% also include a separate encryption key. However, only 20.5% of keys specify the `alg` parameter, with RS256 being the dominant choice. The TA is the only entity using an elliptic curve key, though it appears unused in practice.

Key reuse. While most keys (94.18%) are unique across the federation, a small number are reused extensively: one key appears 1,760 times, and another is shared by an IE and all 7,599 subordinate RPs, suggesting centralized hosting or shared infrastructure. Additionally, about 26.81% of entities reuse their federation keys in their protocol metadata, which is discouraged by the specification. This issue appears in subordinates of eight IEs, others avoid it entirely. The TA is a partial outlier, with 83.5% of its subordinates affected.

Takeaway

Usage of cryptography is solid; however, key reuse within and between entities, omitted `alg` headers, and missing encryption keys can be improved.

5.6. STRIDE-Based Risk Assessment of Issues

We perform a structured STRIDE-based risk assessment [28] to systematically categorize the security implications of our empirically identified issues. Unlike prospective threat modeling, we apply STRIDE retrospectively to analyze the threat categories associated with observed misconfigurations. Our modeling is based on three kinds of attackers: Malicious IEs, malicious LEs, and an external attacker. Details on their capabilities, potential attack targets, and exploitable conformance issues are provided in Table 5 in Appendix D. The mapping between findings and STRIDE components is shown in Table 1.

Table 1. STRIDE-BASED MAPPING OF CONFORMANCE ISSUES.

ID	Finding	STRIDE Mapping
(1)	Key management issues	(S) (T) (R) (E)
(2)	Topology inconsistencies	(T) (D) (E)
(3)	Misconfigured or ineffective policy	(T) (E)
(4)	Trust Mark misconfigurations	(S) (D)
(5)	Excessive statement/TM lifetimes	(S)
(6)	Unavailable / inconsistent SSs	(D)
(7)	Version drift and mixed semantics	(D)

Key management issues. We observe instances of key reuse and cases where the same key is used both as a federation-signing and as an OIDC protocol key. This significantly weakens the authenticity and integrity guarantees of the federation, as compromise of any single holder of a shared key enables an attacker to impersonate all entities using that key, to mint or alter ESs, and in some cases to escalate from protocol-level to federation-level capabilities. These conditions represent clear *Spoofing*, *Tampering*, *Elevation of Privilege*, and *Repudiation* threats. Indirectly, key sharing introduces secondary threats, as compromised keys can be abused to induce availability failures and to receive and decode data intended for impersonated entities. *Information Disclosure* and *Denial of Service* threats only arise as outcomes of key compromise.

Topology inconsistencies and multi-path trust. The current CIE deployment exhibits multiple permissible trust paths for the same entity, including cases where RPs are simultaneously subordinate to an IE and the TA. These inconsistencies in theory allow bypassing intermediate policy enforcement, thereby weakening or nullifying constraints that a superior entity intended to impose. Such ambiguity constitutes *Tampering* in the sense of compromising the integrity of the policy-application process. They may also result in *Elevation of Privilege* when entities can operate under laxer constraints simply by relying on an alternative trust path. Furthermore, inconsistent topology information can cause chain-construction failures, leading to limited but real *Denial of Service* threats.

Misconfigured metadata policies. Metadata policies expressed in SSs are intended to enforce security requirements on subordinate entities. In practice, we find malformed, misplaced, or ineffective policies, which can cause verifiers to silently ignore constraints that were meant to restrict subordinates. This degrades the integrity of the policy mechanism, constituting *Tampering*, and enables *Elevation of Privilege* when subordinates operate with more permissive or insecure metadata than intended.

Trust Mark misconfigurations. Although TMs are usually well-formed and correctly scoped, we identify a small number of mis-issued or structurally irregular TMs. Incorrectly scoped or malformed TMs can cause verifiers to misinterpret an entity's compliance status, enabling limited forms of *Spoofing* when an entity holds a TM not actually issued to it. TM parsing failures may also cause verifiers to reject otherwise valid statements, producing localized *Denial of Service*. TMs can theoretically influence attribute-release decisions. While the current CIE deployment does not appear to use TMs for this purpose, such TM errors could result in *Information Disclosure* threats in the future.

Excessive statement lifetimes. While ESs and TMs exhibit very long lifetimes in some cases, entities do not publish historical keys in their metadata. These conditions increase the window in which compromised keys or stale statements remain valid, directly amplifying *Spoofing* threats by allowing attackers to reuse old yet still-accepted artifacts. The absence of historical keys also complicates post-hoc validation and accountability, which weakens *Repudiation*-related assurances. However, this effect is secondary to the primary authenticity risk.

Unavailable subordinate statements. We were unable to construct some trust chains because the SS endpoints were unreachable or returned errors. These cases do not create integrity or authenticity vulnerabilities, nor do they expand an attacker's capabilities. Instead, they represent purely operational fragility that manifests as *Denial of Service*, since verifiers are unable to construct complete trust chains for the affected entities.

Specification version drift. The federation nominally adheres to OIDFed draft version 24, yet various entities exhibit a mixture of claim sets, field names, and processing rules from later specification drafts. Such semantic divergence can cause verifiers that implement newer drafts to reject statements produced under older interpretations, and vice versa, which may result in interoperability failures. These discrepancies therefore constitute a form of *Denial of Service* threat, but do not meaningfully affect authenticity, integrity, confidentiality, or privilege boundaries.

Takeaway

Our analysis reveals that the identified conformance issues primarily map to Denial of Service threats. Tampering, Elevation of Privilege, and Spoofing threats are also present and almost as common. Notably, we found no conformance issues that enable direct Information Disclosure attacks against other entities. Key management issues appear to be most severe, as they map to four STRIDE components at once.

6. Discussion

In the following, we discuss how our results relate to the research questions and highlight several aspects that warrant further interpretation. We also provide actionable recommendations concerning implementers and operators of OIDFed-based infrastructures, as well as ideas for improving the draft itself and for future research.

6.1. RQ1: General Prevalence

Our results show that the Italian federation adopts many of the core characteristics defined by the OIDFed specification. All ESs correctly implement required JWT claims and almost all entities include valid FE metadata. Informational metadata, such as contacts and homepages, is also widely adopted, reflecting a strong emphasis on discoverability and transparency. Additionally, TMs are extensively used in both ECs and SSs.

However, several specification-defined features are absent or unused. Notably, no entities include the `crit` or `policy_language_crit` claims, meaning that no extension mechanisms to metadata and policy definitions requiring mandatory processing are used. This may reflect either a lack of custom policy extensions or reluctance to risk interoperability issues in a live deployment. The `historical_keys_endpoint`, which supports long-term validation of trust chains across key rotations, is not implemented by the TA. This could hinder the federation's ability to verify past statements once the TA's keys change. Finally, naming constraints are not employed, even though domain name patterns suggest they could help enforce organizational boundaries. In summary, the CIE federation implements the most crucial and widely-used features of OIDFed. However, advanced extensions and long-term key management features remain unused, likely reflecting the federation's short history and focus on stability.

6.2. RQ2: Policy Features

Advanced policy-related features are actively used across the federation, though their enforcement and correctness vary by issuer and claim type. In 39.31% of SSs, the `metadata_policy` claim appears, mostly targeting the RP entity type and defining constraints on core OIDC parameters. This is attempted as well in the 79 cases mentioned in the SS analysis. However, due to the misplacement or misspelling of the claim, the intended policies should be ignored in the trust chain resolution process, rendering them technically non-enforceable. Even though these cases are only considered as warnings, such misconfigurations could be impactful. In the Italian federation, we did not observe any negative consequences during manual verification of the resolved metadata. We found that the respective entities' metadata included the comprehensive set of policies delegated by the TA, while the misplaced or misspelled claims were simply ignored.

At the federation root, the TA issues a single policy targeting the OP and defining strict constraints on supported algorithms, endpoints, scopes, and client registration methods. Many claims are marked as essential, and the OP's metadata mirrors this policy closely, suggesting strong alignment and enforcement. Interestingly, all SSs the TA issues to IEs include a comprehensive set of policies directed at RPs, closely resembling the SSs those IEs issue to their subordinate RPs. This indicates that policy enforcement across the federation is largely driven in a top-down manner by the TA. Additionally, the widespread inclusion of `jwks` in metadata policies suggests that many entities rely on their superiors for key material management. Together, these findings underscore the federation's strongly centralized operational model.

The `constraints` claim appears in only 10.93% of SSs and is rarely used meaningfully. While 97.88% consist of empty objects, these were issued by only two IEs and carry no effective restrictions. However, the remaining 22 SSs that define meaningful constraints were all issued by the TA and target IEs, enforcing that leaf entity types must be RPs. The TA also enforces a global depth restriction via a `max_path_length` constraint in its EC. These results mirror findings from metadata policies, as meaningful constraints are only set by the TA in a top-down approach. Lastly, the federation lacks naming constraints entirely, although many entities share similar domain patterns.

TMs are widely and consistently used, as prescribed in the CIE documentation. Nearly all entities include matching TMs in both ECs and SSs, and most have valid signatures. We identified minor inconsistencies: some entities presented TMs not issued to them, and one EC included a TM with an invalid claim structure. Despite these edge cases, TM usage is well-aligned with the federation's requirements. Almost all entities provide exactly one TM, and almost all use one of the TA-defined TM IDs. All IEs correctly scope the TMs they issue, and no RP uses a TM ID reserved for IEs. However, eight IEs do not issue any TMs despite being authorized to do so, and the only TMI in the federation does not issue any TMs either. Nevertheless, TM distribution remains consistent and policy-compliant in almost all cases, reinforcing their role in trust signaling across the federation.

6.3. RQ3: Conformance

The CIE federation mostly conforms to the OIDFed specification draft. Parts of the federation have evolved beyond draft version 24, while others have not. Therefore, the results of conformance analyses differ greatly, depending on the version used as a basis. When evaluated against the finalized OIDFed 1.0 specification, virtually no entity fulfills all requirements. In particular, claims such as `authority_hints` are now explicitly disallowed in SSs. TMs must also provide the `trust_mark_type` claim instead of the `trust_mark_id` or `id` claim used in former draft versions. These shifts render large parts of the current federation technically non-compliant, underscoring the challenges of deploying against an evolving standard. However, this does not indicate flawed implementations, but rather reflects the evolving nature of the specification, which has undergone substantial changes. The remainder of our discussion is therefore based on conformance analysis tailored to draft version 24 only.

Since most errors and warnings originate from a small subset of IEs, the overall conformance of the federation is disproportionately impacted. Many operators ensure full or almost full conformance within their domain of responsibility. This indicates that targeted remediation could significantly improve the overall compliance, without requiring large-scale structural changes.

Our analysis of `iat` claims shows that almost all entities issue their ESs at the time of request. This can indicate that no caching is used, or that entities are rarely accessed, which is improbable given the vast user base. Some ECs claim to be issued after the request, indicating minor server-side time synchronization issues. The maximum `exp` value we observed was set to more than 1,900 years in the future, indicating incorrect, manual configuration. A majority of entities, most of them related to one specific IE, issue ESs with a lifetime of about 20 years. While the specification does not prescribe specific lifetimes, such extended validity periods appear misaligned with the dynamic and flexible operation that OIDFed is designed to support. The remainder mostly uses lifetimes of 30 minutes, one day, two days, and one year, which can be considered adequate. These shorter lifespans can help ensure compliance of federation members, especially in contexts with evolving security or regulatory requirements. As such frequent renewal is used in practice, long-lived statements are a result of configuration choices rather than technical limitations. Finally, the high similarity between EC and SS lifetimes indicates that organizations deploy similar configurations for their IEs and LEs.

Our analysis reveals almost no issues related to cryptography, since all ESs were correctly signed using the recommended RS256 algorithm. While a single elliptic curve signing key was part of the TA's metadata, no associated signatures were found, suggesting preparation for future use or test deployment. However, we found notable key reuse across the federation: In some ECs, the same key material is used for both federation-level trust and protocol-level operations, despite the specification explicitly discouraging this practice. We also observed identical keys appearing in multiple distinct entities. While such an approach may simplify deployment or key management, it diverges from widely accepted cryptographic best practices, which recommend key separation by function and individual keys for each entity. As argued in the STRIDE assessment, sharing the same key across roles and entities increases the potential impact of any key compromise, as a single leaked key could be utilized to impersonate multiple entities. These findings suggest a need for improved operational guidance and enhanced tooling to support secure key management in large-scale OIDFed deployments.

6.4. RQ4: Topology and Evolution

While the CIE federation saw a dramatic increase in size over our scanning period, the topology remained the same, with a clear structure consisting of a single root of trust, multiple subordinate organizations, and thousands of individual services. The number of IEs increased fivefold, and the total entity count increased by a factor of seven. Most notably, the addition of an IE related to the Italian Ministry of Education and Merit added over 7,000 RPs to the federation. This indicates that the migration of the existing service landscape, as well as the integration of new organizations and RPs was successful. Therefore, OIDFed appears to meet the scaling requirements for large-scale, national trust infrastructures.

However, we observed numerous topological inconsistencies in subordinate-superior relationships: We identified several cases in which entities are considered to be immediate subordinates of two different superiors. Furthermore, we found entities that consider both their IE and the TA as their immediate authorities. This results in a non-tree-shaped federation topology and introduces multiple valid trust paths between a single entity and the TA. Such configurations can lead to ambiguities in trust chain

resolution which may result in incorrect policy application. When intermediates define policies or metadata, those may be bypassed entirely if a trust chain is constructed directly via the TA. While such relationships are not prohibited by the specification draft version used in the deployment, it raises important concerns about policy enforcement and metadata consistency. Notably, from version 41 onward, such scenarios are explicitly highlighted and discouraged in the implementation considerations of the draft.

In addition, some IEs issued SSs to entities that were not present in their own subordinate listing. This discrepancy suggests a potential implementation error at the IE, contradicting the intended semantics of the trust model. Such inconsistencies render the federation's topology ambiguous and may lead to errors in trust chain validation or break assumptions made by RPs.

6.5. Security Implications

Our STRIDE-based analysis shows that only a subset of the observed irregularities translate into substantive authenticity or integrity risks. Key management, topology inconsistencies, metadata-policy misconfigurations, and TM issues stand out as the primary drivers of these risks. These weaknesses expose the federation to impersonation and policy-bypass scenarios if a key is ever compromised or if misconfigurations propagate. However, the analysis also highlights that availability issues, though often caused by benign operations such as specification version mismatches or unreachable endpoints, should not be neglected. In a national digital identity system that underpins access to healthcare, taxation, welfare, and other essential public and private services, even localized or intermittent unavailability can have direct societal impact. For end users, such failures do not appear as abstract trust chain errors but as the inability to authenticate to services they depend on, often at time-critical moments.

While most of the availability shortcomings we observe are not attacker-induced, their presence indicates operational fragility that, if left unaddressed, could become a single point of failure for critical services. Thus, the practical security implications of the issues we identify extend beyond classical adversarial threat models and underscore the need for stronger operational resilience and governance mechanisms in large-scale identity federations.

6.6. Implications Beyond Italy

Our analysis shows that nearly all policies within the CIE federation are delegated by the TA, with IEs rarely defining their own subordinate constraints. This top-down enforcement model aligns with the centralized nature of Italy's national deployment, where the TA serves as both technical and governance authority. However, in broader and more decentralized ecosystems, such as the anticipated EUDI Wallet infrastructure or the academic inter-federation eduGAIN, a purely top-down model may not be optimal. In that context, a multi-anchor or hierarchical trust structure may emerge, where a pan-European TA could enforce cross-border baseline policies, e.g., conformance to the General Data Protection Regulation. National IEs could then enforce country-specific requirements such as assurance levels, credential formats, or sectoral regulations. This layered approach would allow for regulatory subsidiarity while preserving interoperability across the federation. Supporting such flexibility through the `metadata_policy` and `constraints` mechanisms will be critical for cross-national deployments. Additionally, extensibility through the `crit` and `policy_language_crit` claims would likely be required to enable context-specific enforcement in federations spanning multiple legal and policy domains. We therefore treat CIE as a first-of-kind case study rather than a representative sample of future OIDFed deployments: the observed operational choices are Italian, but the underspecified behaviors they expose are relevant to any deployment relying on the same draft semantics.

6.7. Actionable Recommendations

OIDFed is expected to be deployed in more high-profile infrastructures in the coming years. In addition, stakeholders within the R&E sector are in the process of transitioning their existing, SAML 2.0-based service landscape to OIDFed [39, 35]. Therefore, we put forward the following actionable recommendations based on our results. Specifically, we consider the specification itself, software developers, current and future federation operators, and directions for future research.

Specification. Other OpenID specifications like the Shared Signals Framework [38] include a `spec_version` parameter. The provision of such versioning information in the entity metadata would improve interoperability. While this would be especially helpful during the draft stage, updates to the specification after its initial release are also possible. The specification already recommends only constructing tree topologies to prevent loops and remove ambiguity during trust resolution. The multilateral nature of OIDFed makes it impossible to enforce these rules. However, these recommendations could be further emphasized to guide implementers. Finally, the specification could recommend sensible default lifetimes, as we observed needlessly long time periods in many instances.

Software Developers. Implementers should test their solutions for interoperability with other implementations. Specifically, it is important that implementations targeting different draft versions are compatible or at least fail gracefully in case of errors. During our analysis, we found some configuration issues that could have been avoided through strict and comprehensive validation. Software developers should aim to implement such checks to ensure correct operation.

Federation Operators. Throughout a federation, operators should make sure that all deployed software supports a single version of the OIDFed specification. This reduces interoperability risks and allows for automated compliance checks. Such compliance validation should be done at regular intervals to spot issues early on. Operators should also monitor the topology of their federation, ensuring congruence between subordinate listings and authority hints for all entities.

Future Research. Since OIDFed 1.0 has now been standardized, an immediate research question is how existing draft-based deployments migrate to the final specification without breaking deployed clients and RPs. The introduction of OIDFed in the vast and heterogeneous

R&E sector could surface new challenges and limitations. Assessing these developments through studies similar to ours would enable comparisons between approaches to implementation and operation. Furthermore, qualitative research about the experiences of users, implementers, and operators with OIDFed could complement our quantitative analyses. Finally, we have developed a web-based dashboard for monitoring these developments. Directions for accessing our public instance are provided at publication time in our supplemental repository [30].

7. Conclusion

Our comprehensive 15-month analysis of the Italian CIE federation demonstrates the viability and operational stability of the OIDFed framework in a national production deployment. The Italian ecosystem grew from 1,584 to 12,789 entities, a sevenfold expansion largely driven by the onboarding of Ministry of Education and Merit services. Despite this rapid growth, almost all published metadata was correctly signed, and the employed cryptography was implemented without notable issues, underscoring the robustness of OIDFed's core design.

Beyond this positive picture, our study identifies several structural weaknesses that affect the security posture of the federation. Through our STRIDE-informed assessment, we find that only a small subset of deviations translates into meaningful security risks. Specifically, we detected extensive key material reuse, topology inconsistencies that could enable policy bypass in certain trust paths, and misconfigured or ineffective metadata policies. These issues do not indicate ongoing exploitation, but they substantially amplify the consequences of any future key compromise or malicious insider, and they undermine the integrity of delegated governance. At the same time, irregularities such as misconfigured or expired TMs, unavailable SSs, and specification drift highlight operational fragilities that can cause availability failures. In a national digital identity system classified as critical infrastructure, even localized unavailability can directly translate into citizens being unable to authenticate to essential public and private services, elevating availability to a major security concern.

Most other conformance issues we observed such as misspelled claims, non-standard headers, or inconsistent JWKS labels have limited direct security impact and are best understood as robustness or maintainability concerns. Nonetheless, they collectively point to a need for improved governance and validation across the federation.

We therefore see three immediate priorities to address these shortcomings. First, specification authors should formalize version signaling and encourage single-version rollouts to reduce the draft drift we observed. Second, implementers should adopt stronger validation tooling to catch key reuse and conformance pitfalls before deployment. Third, operators would benefit from monitoring solutions capable of surfacing topology inconsistencies and policy deviations in real time, improving both security and operational resilience. Our artifacts can serve as a foundation for such a monitoring dashboard.

These findings matter well beyond Italy. As the forthcoming EUDI Wallet will require a trust layer that can scale easily while enforcing fine-grained policy in an interoperable way, our analysis suggests that OIDFed is a credible candidate. Likewise, R&E ecosystems like eduGAIN and national federations can leverage OIDFed to modernize their largely static SAML 2.0 deployments. Italy's OIDFed journey confirms that this protocol has the potential to shoulder nation-scale digital identity systems that serve as critical infrastructure for public and private services. The broader contribution of this study is therefore not only the characterization of CIE, but also a reusable framework for detecting where OIDFed specifications, implementations, and federation operations interact in security-relevant ways. By following the lessons given in this study, the protocol and its implementations can be further refined to ultimately deliver a safer and more interoperable backbone of digital identity services across Europe and beyond.

Ethics Considerations

For our ethical assessment, we primarily considered the guidelines laid out in The Menlo Report [3]. Before we began our scans, we consulted with one of the project leaders responsible for developing the CIE federation and received approval for conducting this study. In addition, we self-assessed our research methodology via a form provided by the ethics committee of our institution. Based on this, no motion by the committee was required.

All data were stored on a Virtual Machine hosted within our faculty network, accessible only to the authors of this study, to ensure data security and privacy. The 24-hour scanning interval with a single request per entity, and an audience of nearly 60 million people in Italy, resulted in negligible additional load on live services and infrastructure.

Although all data we analyzed in this study was obtained from publicly available sources on the Internet, we considered potential harm that could arise from this research. The `contact` claim, part of an entity's `federation_entity` metadata, can contain email addresses. However, such addresses serve an operational purpose and are explicitly published as public contact points. We found no data that can be considered personally identifiable information in our dataset.

We also considered potential harm that could arise from aggregated data. Therefore, we established a protocol on how to handle security and conformity issues prior to publication. In case of major issues that could entail security risks for operators or users, we decided to conduct a responsible disclosure process. As no major security issues were apparent, such a procedure was not necessary. However, we contacted one IE operator via email on 2025-04-04 to report issues with the `cty` header parameter. In addition, we notified 18 IE operators on 2025-04-30 about all conformance issues affecting their subordinates and provided them with further preliminary analysis results. We also shared all of our results with the aforementioned project leader who forwarded our conformance issue reports to IE operators that only accept certified email, i.e., Posta Elettronica Certificata. Finally, we assessed our methodology to be fully compliant with applicable law.

Data Availability

We provide the complete dataset, together with our tooling, as an artifact [30].

Disclosure of Interests

The authors have no competing interests to declare that are relevant to the content of this article.

Acknowledgment

We thank Giuseppe De Marco for discussing our initial idea and for helpful reviews. We also thank Diana Gudu for the development of ofcli. The authors employed Claude [2] and ChatGPT [44] to improve spelling and grammar.

References

[1] Agenzia per l'Italia Digitale. *Three-Year ICT Plan — National platforms that provide services to citizens/businesses or other public administrations*. 2025. URL: https://docs.italia.it/italia/piano-triennale-ict/pianotriennale-ict-doc/it/2024-2026/capitolo-4_piattaforme/piattaforme-nazionali-che-erogano-servizi-a-cittadiniimprese-o-ad-altre-pa.html (visited on 11/20/2025).

[2] Anthropic. *Claude Sonnet 4.5*. https://claude.ai. Version claude-sonnet-4-5-20250929. Large language model and AI assistant. 2025.

[3] Michael Bailey et al. "The Menlo Report". In: *IEEE Security & Privacy* 10.2 (2012), pp. 71–75. DOI: 10.1109/MSP.2012.52.

[4] Scott O. Bradner. *Key words for use in RFCs to Indicate Requirement Levels*. RFC 2119. Mar. 1997. DOI: 10.17487/RFC2119. URL: https://www.rfc-editor.org/info/rfc2119 (visited on 11/20/2025).

[5] CIE Federation. *OpenID Federation EID Intermediate Entity*. 2025. URL: https://eid.istruzione.it/eid-gateway-federation-services (visited on 11/20/2025).

[6] CIE Federation. *OpenID Federation OpenID Provider*. 2025. URL: https://oidc.idserver.servizicie.interno.gov.it/ (visited on 11/20/2025).

[7] CIE Federation. *OpenID Federation Trust Anchor*. 2025. URL: https://oidc.registry.servizicie.interno.gov.it (visited on 11/20/2025).

[8] European Commission. *European Digital Identity Wallet*. 2025. URL: https://eu-digital-identity-wallet.github.io/eudi-doc-architecture-and-reference-framework/latest/ (visited on 11/20/2025).

[9] Connect2id. *OpenID Federation 1.0 and the trust chain explained*. 2024. URL: https://connect2id.com/learn/openid-federation (visited on 11/20/2025).

[10] Cristiano Ghidotti - Punto Informatico. *Documents on IO: 5 million activations for IT-Wallet*. 2025. URL: https://www.punto-informatico.it/documenti-su-io-5-milioni-attivazioni-itwallet/ (visited on 11/20/2025).

[11] Cristiano Ghidotti - Punto Informatico. *IT-Wallet System*. 2025. URL: https://innovazione.gov.it/progetti/sistema-it-wallet/ (visited on 11/20/2025).

[12] Giuseppe De Marco. *OpenID Federation 1.0 and EUDIW TL / X.509 PKI in IT-Wallet*. Presentation at the 4th International Workshop on Trends in Digital Identity, Verona. Accessed 30 April 2026. 2026. URL: https://peppelinux.github.io/Wallet-Presentations/openid-federation-wallet-tdi/index.html.

[13] Giuseppe De Marco. *Personal communication*. Apr. 2025.

[14] Giuseppe De Marco and Francesco Antonio Marino. *The Italian National Federation: Design, Deployment, and Lessons Learned*. 2024. URL: https://www.kuppingercole.com/watch/italian-national-federation-eic24 (visited on 11/20/2025).

[15] Giuseppe De Marco et al. *OpenID Federation Wallet Architectures 1.0 - draft 03*. The OpenID Foundation. 2024.

[16] Ministero dell'Interno. *Carta di Identità Elettronica (CIE) - Funziona, semplicemente*. 2025. URL: https://www.cartaidentita.interno.gov.it/ (visited on 11/20/2025).

[17] Department for Digital Transformation, IPZS Istituto Poligrafico e Zecca dello Stato, PagoPA, AGID Agency for Digital Italy. *GitHub - italia/eid-wallet-it-docs*. 2025. URL: https://github.com/italia/eid-wallet-it-docs/blob/versione-corrente/docs/en/trust.rst (visited on 11/20/2025).

[18] Department for Digital Transformation Italy. *GitHub - openid-federation-browser*. 2025. URL: https://github.com/italia/openid-federation-browser (visited on 11/20/2025).

[19] Agenzia per l'Italia Digitale. *SPID - Sistema Pubblico di identità Digitale*. 2025. URL: https://www.spid.gov.it/ (visited on 11/20/2025).

[20] Agenzia per l'Italia Digitale. *SPID/CIE OIDC - Regole Tecniche version: latest documentation*. 2025. URL: https://italia.github.io/spid-cie-oidc-docs/en/ (visited on 11/20/2025).

[21] eduGAIN and GÉANT. *What is eduGAIN - eduGAIN*. 2025. URL: https://edugain.org/about-edugain/what-is-edugain/ (visited on 11/20/2025).

[22] DC4EU: Digital Credentials for Europe. *DC4EU Feedback: OpenID Federation and dPKI Trust Frameworks support*. 2025. URL: https://github.com/eu-digital-identity-wallet/eudi-doc-architecture-and-reference-framework/issues/224 (visited on 11/20/2025).

[23] OpenID Foundation. *R&E Working Group – Charter*. 2025. URL: https://openid.net/wg/rande/charter/ (visited on 11/20/2025).

[24] OpenID Foundation. *The OpenID Federation Interoperability Event*. 2025. URL: https://openid.net/the-openid-federation-interoperability-event/ (visited on 11/20/2025).

[25] gematik. *Spezifikation Federation Master 1.4.1*. 2025. URL: https://gemspec.gematik.de/docs/gemSpec/gemSpec_IDP_FedMaster/latest/ (visited on 11/20/2025).

[26] Diana Gudu. *dianagudu/ofcli: A helper tool for exploring OIDC federations*. 2025. URL: https://github.com/dianagudu/ofcli (visited on 11/20/2025).

[27] Roland Hedberg et al. *OpenID Federation 1.0*. OpenID Final Specification. 2026. URL: https://openid.net/specs/openid-federation-1_0-final.html (visited on 04/20/2026).

[28] Shawn Hernan et al. *Uncover Security Design Flaws Using the STRIDE Approach*. MSDN Magazine. 2006. URL: https://msdn.microsoft.com/en-us/magazine/cc163519.aspx (visited on 11/18/2025).

[29] Paul den Hertog, Niels van Dijk, and Klaas Wierenga. *A Trusted Foundation for the EUDI Wallet in Research and Education: Why eduGAIN and OpenID Federation Matter*. 2025. URL: https://gist.github.com/pauldenhertog/5bebc3992e0ce831122ac3f37c7936ad (visited on 11/20/2025).

[30] Tobias Hilbig, Erwin Kupris, and Thomas Schreck. *Artifact for "All You Need is Trust: A Longitudinal Analysis of Italy's OpenID Federation Journey"*. 2026. URL: https://github.com/hm-seclab/paper-openidfed-italy (visited on 05/01/2026).

[31] Michael B. Jones, John Bradley, and Nat Sakimura. *JSON Web Token (JWT)*. RFC 7519. May 2015. DOI: 10.17487/RFC7519. URL: https://www.rfc-editor.org/info/rfc7519 (visited on 11/20/2025).

[32] Michael B. Jones and Nat Sakimura. *JSON Web Key (JWK) Thumbprint*. RFC 7638. Sept. 2015. DOI: 10.17487/RFC7638. URL: https://www.rfc-editor.org/info/rfc7638 (visited on 11/20/2025).

[33] Erwin Kupris et al. "A-WAYF: Automated Where Are You From in Multilateral Federations". In: *2nd International Workshop on Trends in Digital Identity*. CEUR Workshop Proceedings, 2024, pp. 6–17. URL: http://ceur-ws.org/Vol-3863/paper1.pdf.

[34] T. Lodderstedt, K. Yasuda, and T. Looker. *OpenID for Verifiable Credential Issuance - draft 16*. The OpenID Foundation. 2025. URL: https://openid.net/specs/openid-4-verifiable-credential-issuance-1_0-16.html (visited on 11/20/2025).

[35] Giuseppe De Marco et al. *OpenID for Research and Education (OpenID4RE) Project - GitHub*. 2025. URL: https://github.com/GEANT/edugain-openidfed (visited on 11/20/2025).

[36] Francesco Antonio Marino and Pasquale Cerqua. *Trust Frameworks in Digital Identity - Building Bridges Between EUDIW and OpenID Federation*. 2025. URL: https://st.fbk.eu/assets/areas/events/TDI2025/slides/2_1_Marino_Cerqua.pdf (visited on 11/20/2025).

[37] OASIS. *Security assertion markup language (SAML) V2.0*. 2008. URL: https://docs.oasis-open.org/security/saml/Post2.0/sstc-saml-tech-overview-2.0.html (visited on 11/20/2025).

[38] OpenID Foundation. *Shared Signals - A Secure Webhooks Framework — OpenID*. 2017. URL: https://openid.net/wg/sharedsignals/ (visited on 11/20/2025).

[39] Wolfgang Pempe. *OpenID Connect Federation - 80. DFN-Betriebstagung*. 2024. URL: https://www.dfn.de/wp-content/uploads/2023/10/openid_federation.pdf (visited on 11/20/2025).

[40] Daniela Pöhn and Wolfgang Hommel. "An overview of limitations and approaches in identity management". In: *Proceedings of the 15th International Conference on Availability, Reliability and Security*. ARES '20. Virtual Event, Ireland: Association for Computing Machinery, 2020. ISBN: 9781450388337. DOI: 10.1145/3407023.3407026.

[41] *Regulation (EU) 2024/1183 of the European Parliament and of the Council of 11 April 2024 amending Regulation (EU) 910/2014 as regards establishing the European Digital Identity Framework*. Apr. 2024.

[42] *Regulation (EU) 910/2014 of the European Parliament and of the Council of 23. July 2014 on electronic identification and trust services for electronic transactions in the internal market and repealing Directive 1999/93/EC*. Sept. 2014.

[43] Natsuhiko Sakimura et al. *OpenID Connect Core 1.0*. The OpenID Foundation. 2014.

[44] Aaditya Singh et al. *Introducing GPT-5 — OpenAI*. 2025. URL: https://openai.com/en-US/index/introducing-gpt-5/ (visited on 11/18/2025).

[45] SPID Federation. *OpenID Federation Trust Anchor*. 2025. URL: https://registry.spid.gov.it (visited on 11/20/2025).

[46] O. Terbu et al. *OpenID for Verifiable Presentations - draft 25*. The OpenID Foundation. 2025.

Appendix A.
OIDFed Field Semantics

Table 2. SECURITY-RELEVANT OPENID FEDERATION FIELDS REFERENCED IN OUR ANALYSIS. EXPLANATIONS FOLLOW OIDFED 1.0 TERMINOLOGY.

Field	Location / object	Purpose in OIDFed	Why misuse or misconfiguration matters
`alg`	JWS header of Entity Statements, Trust Marks, and related JWTs	Identifies the JWS signing algorithm used to protect the JWT; OIDFed 1.0 requires an acceptable signing algorithm and forbids `none`.	An absent, unsupported, or weak algorithm can cause statement rejection or undermine signature-validation assumptions.
`kid`	JWS header; matched against the issuer's `jwks`	Identifies the public key that verifiers should use to validate the JWT signature; the value must match a key identifier in the issuer's JWKS.	Incorrect or reused key identifiers can make validation fail, select the wrong key, or obscure key-rotation and key-compromise boundaries.
`iat`	JWT claim in Entity Statements and Trust Marks	States when the object was issued, expressed as seconds since the Unix epoch; verifiers check that the current time is after this value.	Future or implausible issue times can indicate clock drift or faulty generation logic and may cause valid objects to be rejected.
`exp`	JWT claim in Entity Statements and Trust Marks	States when the object expires; verifiers check that the current time is before this value.	Expired objects cause validation failure, while excessive lifetimes increase the window in which stale statements or compromised keys remain useful.
`authority_hints`	Entity Configuration, except Trust Anchors without superiors	Lists the Entity Identifiers of the entity's Immediate Superiors; it supports discovery of Entity Configurations and construction of candidate trust chains.	Missing, stale, or inconsistent hints can make the federation graph ill-formed, cause discovery failures, or create ambiguous trust-chain resolution behavior.
`metadata_policy`	Subordinate Statement (allowed in Entity Configurations in earlier draft versions)	Defines policy rules that apply to the metadata of the statement subject and its subordinate entities; policies are resolved and applied during trust chain resolution.	Misplaced, malformed, or misspelled policies are not applied as intended, which can silently remove restrictions on subordinate metadata.
`constraints`	Subordinate Statement (allowed in Entity Configurations in earlier draft versions)	Defines trust chain constraints such as path-length limits, naming constraints, or allowed entity types; each present constraint is checked during trust chain resolution.	Incorrect or ineffective constraints can allow unintended topology or entity-type relationships, while invalid constraints can cause otherwise usable chains to be rejected.
`trust_mark_type`	Trust Mark object and Trust Mark JWT	Identifies the type of Trust Mark, i.e., the conformance or accreditation profile being asserted; the value in the containing object must match the value inside the Trust Mark JWT.	A mismatch or deprecated field name can cause Trust Mark validation failure or misrepresent the conformance profile an entity is claiming.
`trust_mark_id`	Legacy / older draft Trust Mark field	Older OIDFed drafts used this field to identify Trust Marks; OIDFed 1.0 uses `trust_mark_type`.	Using legacy field names creates draft-version drift and can break interoperability with OIDFed validators.

Appendix B.
Details from Selected Scans

Table 3. SELECTED SCANNING DATES. L1 AND L2 DENOTE THE NUMBER OF ENTITIES AT THE RESPECTIVE DISTANCE TO THE TA; VALID DENOTES THE SHARE OF ECS WITHOUT WARNINGS OR ERRORS. EVENTS: (1) INCLUSION OF LES FROM DEBUG LOGGING; (2) ADDITION OF EDUCATION-RELATED RPS; (3) INCLUSION OF IES AND LES FROM DEBUG LOGGING; (4) INITIATION OF RESPONSIBLE DISCLOSURE PROCESS.
(*) THE TMI WAS ALMOST CERTAINLY PART OF THE FEDERATION AT THIS TIME, BUT NOT PART OF OUR DATASET DUE TO BEING DROPPED.

Scan date	TAs	OPs	TMIs	IEs	RPs	LEs	Entities	TMs	L1	L2	Valid	Event
2024-05-01	1	1	(1*)	6	1,576	1,577	1,584	1,576	651	932	7.07%	Initial scan date
2024-06-21	1	1	1	5	1,947	1,950	1,956	1,939	942	1,013	12.37%	(1) LE debug logs
2025-03-08	1	1	1	17	10,756	10,760	10,778	10,808	1,793	8,984	78.33%	(2) Education inclusion
2025-03-29	1	1	1	21	12,309	12,313	12,335	12,316	2,592	9,742	68.64%	(3) LE + IE debug logs
2025-04-30	1	1	1	21	12,435	12,438	12,460	12,444	2,639	9,820	68.32%	(4) Responsible disclosure
2025-07-31	1	1	1	21	12,764	12,767	12,789	12,825	2,766	10,022	67.15%	Final scan date

Appendix C.
Dataset Summary Table

Table 4. SUMMARY OF THE LONGITUDINAL DATASET AND FINAL-SCAN SNAPSHOT. UNLESS STATED OTHERWISE, COUNTS REFER TO THE FINAL SCAN ON 2025-07-31.

Longitudinal collection and dataset		**Final-scan topology and conformance**	
Collection		*Topology*	
Scan start	2024-05-01	Entities at distance 1 from TA	2,766
Scan end	2025-07-31	Entities at distance 2 from TA	10,022
Scheduled scan days	456	Maximum observed trust-chain depth	2
Broken scan outputs	48 days (10.52%)		
Final-scan dataset		*EC conformance*	
Entity Configurations (ECs)	12,789	Valid ECs	67.15%
Subordinate Statements (SSs)	12,524	ECs with warnings	20.94%
Distinct Trust Marks (TMs)	12,825	Invalid ECs	11.91%
Entities without constructable trust chains	326 (2.61%)		
Entity types		*SS conformance*	
Trust Anchors (TAs)	1	Valid SSs	75.65%
OpenID Providers (OPs)	1	SSs with warnings	11.69%
Trust Mark Issuers (TMIs)	1	Invalid SSs	12.66%
Intermediate Entities (IEs)	21		
Relying Parties (RPs)	12,764	*Entity Statement conformance overall*	
Leaf Entities (LEs)	12,767	Valid Entity Statements	71.36%
Total entities	12,789	Entity Statements with warnings	16.36%
		Invalid Entity Statements	12.28%

Appendix D.
STRIDE Assessment Details

Table 5. ATTACKER MODEL FOR STRIDE-BASED ANALYSIS.

Attacker Type	**Capabilities**	**Targets**	**Relevant issues**
Malicious IE	• Manipulate own metadata • Access IE/LE private keys • Request SSs • Issue SSs	All subordinates	(1) Key management issues (2) Topology inconsistencies (4) Trust Mark misconfigurations
Malicious RP	• Manipulate own metadata • Request SSs	Sister entities under same IE	(1) Key management issues (2) Topology inconsistencies (4) Trust Mark misconfigurations
External attacker	• Network man-in-the-middle • Request SSs	All entities	(1) Key management issues (2) Topology inconsistencies

Appendix E.
Additional Figures

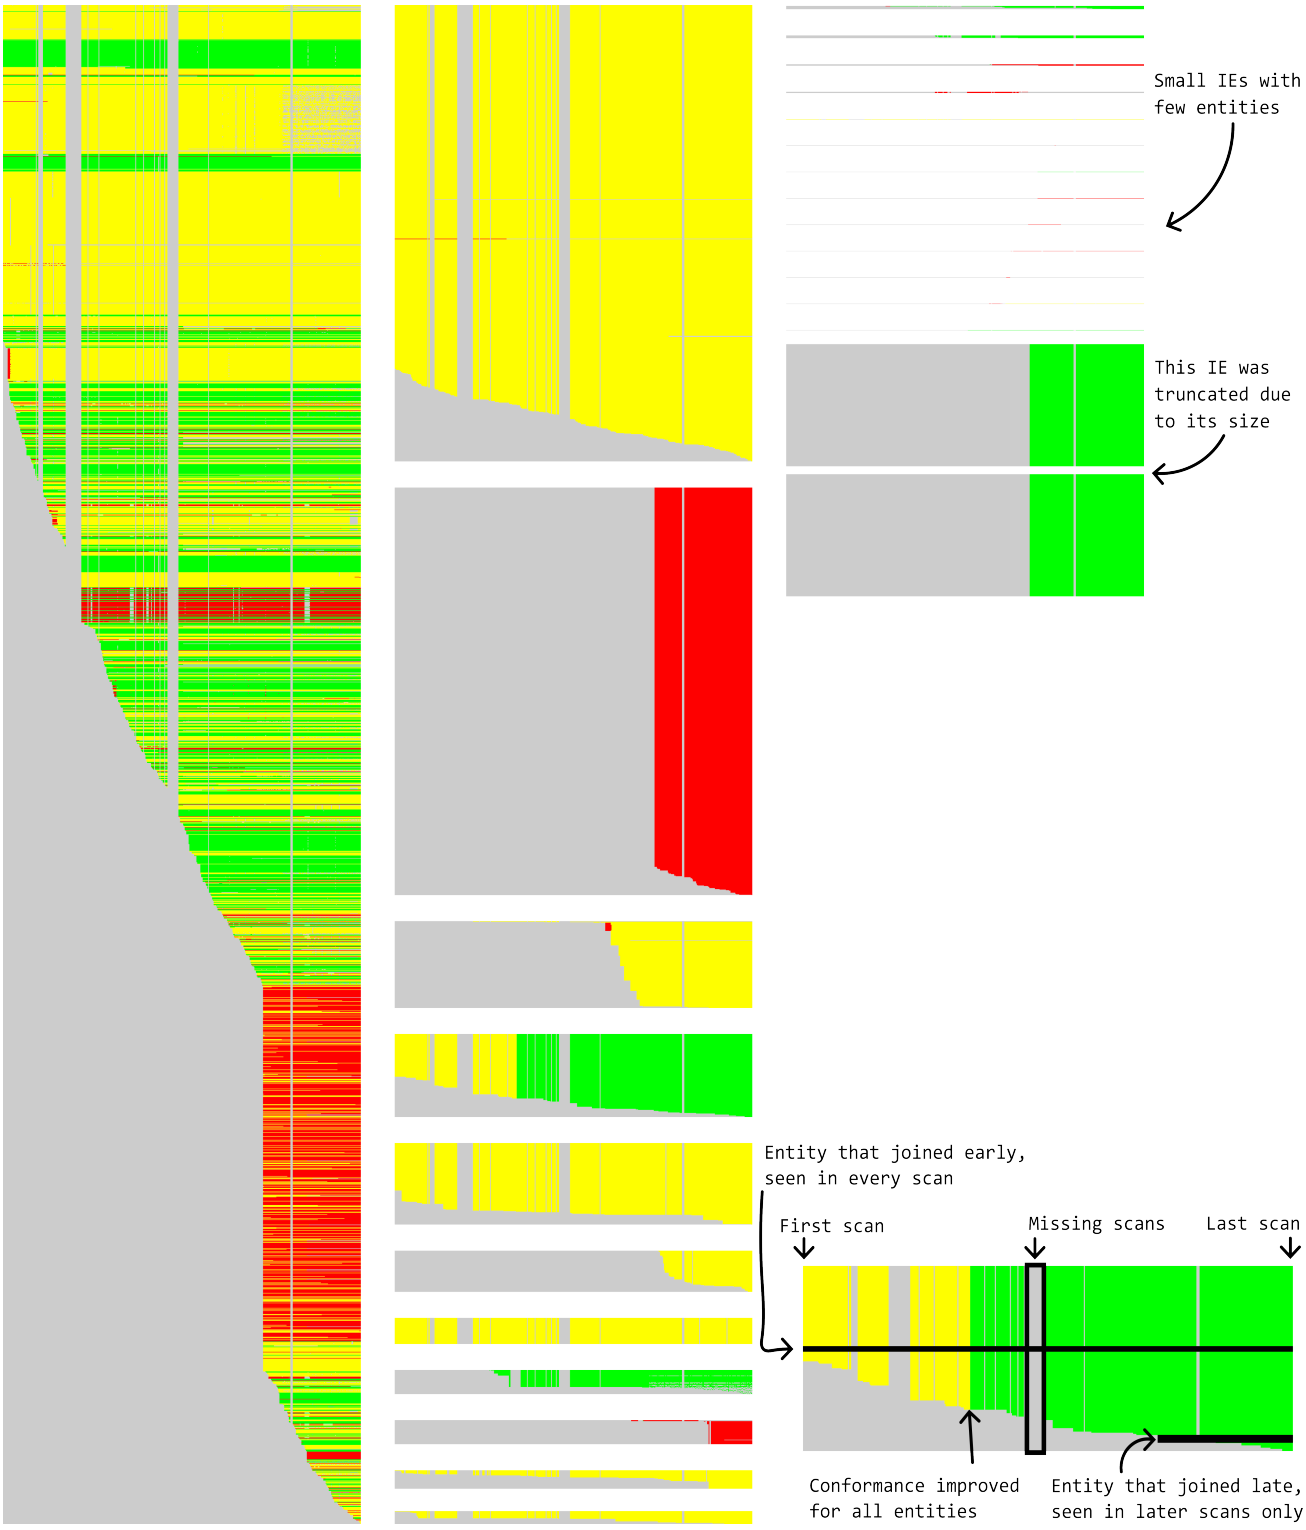

Figure 7. Charts depicting entity validity over time. The x-axis contains all dates between 2024-05-01 and 2025-07-31, including dates at which scans failed. The y-axis contains all entities, grouped by IE, and sorted by the date they first appeared in the dataset. The colors denote each entity's validity at the specific date: green for valid, yellow for warnings, and red for errors. Missing data is depicted in light grey. One specific entity occupies one row, with each column depicting the results from one specific scan. Due to the large number of entities, the chart is split into three parts.