

EU Policies Meet Global Practices: The Discourse on Qualified Website Authentication Certificates

Pratyush Dikshit^a, Pol Hölzmer^a, Johannes Sedlmeir^b, Gilbert Fridgen^a

^a*SnT, University of Luxembourg, Luxembourg, Luxembourg*

^b*Department of Information Systems, University of Münster, Münster, Germany*

Abstract

The revision of the European electronic Identification, Authentication and Trust Services (eIDAS) Regulation (EU) 910/2014 has sparked considerable debate, particularly regarding the mandatory recognition and display of Qualified Website Authentication Certificates (QWACs) by browsers under Article 45. Open letters from security researchers and browser vendors, among others, have highlighted potentially harmful implications for trust, authority, and interoperability, while other stakeholders have labeled the corresponding arguments as disinformation and emphasized anticipated benefits. This study examines the origins, governance context, and arguments surrounding this “QWAC controversy” using a mixed-methods approach. Our data collection relies on a multivocal literature review that encompasses academic publications, position papers, and additional non-academic web documents, as well as semi-structured stakeholder interviews with supporters and critics of the revised Article 45. We map 20 core arguments that emerged from coding of our data into their Toulmin components across three analytical themes – Security & Trust, Governance & Authority, and Compliance & Interoperability – and also derive a corresponding threat model including mitigation levers. The analysis reveals divergent viewpoints, advocacy strategies, and the dissemination of conflicting information amid uncertainty. Indeed, our evaluation demonstrates that many arguments exchanged still hinge on future implementation choices. We clarify the socio-technical trade-offs in EU digital trust policy that collide with the established WebPKI and offer evidence-informed considerations for balancing public and private interests in future discourse and legislative efforts.

Keywords: eIDAS, Extended Validation, Qualified Trust Service Provider (QTSP), Qualified Website Authentication Certificate (QWAC), WebPKI

1. Introduction

The European Union’s (EU) digital single market (DSM) has been undergoing significant transformation in recent years (Barbereau et al., 2024; Heidebrecht, 2024). The introduction of the electronic Identification, Authentication and Trust Services (eIDAS) Regulation (EU 910/2014) marked an important development, as it aimed to enhance identification and authentication-related processes in digital transactions through a set of trust services within the EU (European Commission, 2020b). It affects EU citizens and businesses, as well as non-EU residents, such as guest workers and visitors engaging with services operating in the EU (Weigl et al., 2022).

Qualified Website Authentication Certificates (QWACs) were introduced under Article 45 of eIDAS as a means to establish website authentication as one of these trust services, mandating that users must be able to verify the identity of the organization associated with the website they are interacting with. As such, QWACs not only implement the General Data Protection Regulation’s (GDPR) (EU 2015/2366) mandate to unambiguously identify data processors but also promise to improve accountability in digital interactions, with cross-border recognition across all EU member states (Entschew et al., 2022). QWACs are issued by regulated organizations called qualified trust service providers (QTSPs) and implemented as a type of Transport Layer Security (TLS) certificates (Boeyen et al., 2008), following the baseline requirements (BR) of the CA/Browser Forum (CA/B), which is a consortium mainly consisting of browser vendors and Certification Authorities (CAs) (CA/B Forum, 2025). Although these stakeholders operate outside a specific regulatory framework, they lead the governance of the WebPKI, which represents one of the key foundations of web security (Grindal et al., 2025).

Despite their promises, QWACs have seen sparse adoption since the advent of eIDAS in 2014 (European Commission, 2023b; Entschew et al., 2022). The subsequent discussion to promote QWACs by means of mandatory browser recognition during the revision of eIDAS – known as eIDAS 2.0 Regulation (EU 1183/2024) – has sparked intense controversy among stakeholders, including browser vendors and cybersecurity researchers, who warned of significant risks for the Internet as we know it (Mozilla, 2021, 2023b; Mozilla et al., 2023; Scientists and NGOs, 2023b) on the one side and governmental agencies and QTSPs on the other (ESD, 2022a). However, related scientific work on QWACs is limited and has primarily focused on Article 45 of eIDAS and its legal foundations (Martius et al., 2024; Wazan et al., 2024),

with little attention to the socio-technical controversy across all stakeholders. Our systematic literature review retained only four peer-reviewed studies with a substantial focus on QWAC, and none provides a complete and nuanced view on their integration into existing structures. Several academic and non-academic contributions take a supportive stance, recognizing implementation challenges but emphasizing the potential of QWACs to enhance accountability and trust (Entschew et al., 2022; Schwalm, 2023). However, the interaction between legal provisions, delegated acts, and evolving technical standards with these stakeholder positions remains underexplored. Addressing this gap is essential for a more comprehensive assessment of the opportunities and challenges of QWACs.

We hence examine the QWAC controversy by focusing on the interplay among key stakeholders and the implications of the mandates imposed by Article 45 as modified by eIDAS 2.0. In doing so, we sought to understand the motivations and concerns of each stakeholder group. To offer a balanced analysis, we integrate the potential benefits and drawbacks of QWACs, as expressed by both supporters and critics. Therefore, this study aims to provide a neutral perspective on the controversy following the revision of eIDAS Article 45. We ask: *How does the QWAC controversy surrounding eIDAS Article 45 reshape the relationship between legal mandates and established technical trust infrastructures?*

To answer this research question, we employ a mixed-methods research design. We begin with a systematic literature review to assess the state of peer-reviewed research on QWACs. Given the scarcity of academic studies, we extend the evidence base through a systematic multivocal review of gray literature, covering both normative and prescriptive sources (e.g., regulations, standards, and governance artefacts) and non-academic web documents that capture stakeholder discourse (e.g., position papers, expert essays, professional forum discussions, and news reporting). We analyze all sources through a single, transparent argumentation procedure, while distinguishing discourse capture from normative and technical backing. We then apply Toulmin’s argumentation model (Toulmin, 2003) as a relational coding framework to decompose statements into claims, grounds, warrants, backings, qualifiers, and rebuttals, and iteratively abstract them into a consolidated argument map. Finally, we conduct 15 semi-structured interviews with experts selected using purposive sampling to refine or qualify Toulmin components where the documentary record is weak, ambiguous, or contested, as well as to triangulate how key stakeholders interpret the collected arguments.

Since our analysis particularly yields that many asserted risks and benefits depend on future implementation and enforcement choices, the study qualitatively evaluates security reasoning in the pre-deployment discourse rather than investigating post-deployment outcome metrics.

As such, this study makes four key contributions aligned with Gregor’s contribution types (Gregor, 2006). First, we provide an evidence-grounded reconstruction of the Article 45 controversy in the form of a consolidated stakeholder argument map and a Toulmin-structured argument catalogue (Appendix A.2), organized along three themes: Security & Trust, Governance & Authority, and Compliance & Interoperability. Second, we contribute a transparent, reusable Toulmin-based argument-abstraction procedure for comparing heterogeneous technical, legal, and governance claims by making their underlying assumptions and scope conditions explicit. Third, we offer an explanatory account and governance-impact perspective, arguing that the core tension is governance-of-security, i.e., how legal mandates act as security mechanisms and interact with browser-led WebPKI governance (trust anchors, revocation authority, accountability, and interoperability), rather than the cryptographic soundness of TLS certificates. Fourth, we translate recurring controversy claims into an argument-linked threat model with non-prescriptive mitigation levers (Appendix C), mapping asserted risks to concrete failure modes, affected layers, and key implementation choices (including 1-QWAC vs. 2-QWAC variants) to support pre-deployment trade-off reasoning for implementing eIDAS Article 45.

The remainder of this paper is organized as follows. Section 2 outlines the context of the QWAC controversy, including its regulatory and technical foundations. Section 3 details our qualitative mixed-methods design, including the multivocal literature review, argument coding and abstraction, as well as expert interviews. Section 4 presents the analytical framework, including the stakeholder typology, analytical themes, and shared premises. Section 5 outlines the consolidated stakeholder argument map and synthesis across the three argument classes. Section 6 derives theoretical and practical implications, provides actionable recommendations, and develops an argument-linked threat model. Finally, Section 7 discusses limitations and future work and concludes the paper.

2. Background

This section discusses the three fundamental components of the QWAC controversy: (1) the technical foundation provided by the WebPKI and TLS; (2) the governance framework established by the eIDAS regulation; and (3) the new European Telecommunications Standards Institute (ETSI) technical specification for QWAC certificate profiles, synthesizing the European Commission’s (EC) and CA/Browser Forum’s (CA/B) requirements.

2.1. WebPKI

The public key infrastructure (PKI) of the Web (hereafter WebPKI) represents the backbone of the TLS protocol, securing *virtually all* web traffic (Google, 2025) by providing data confidentiality and endpoint authentication (Rescorla, 2018). In the WebPKI, a root or intermediate CA binds a cryptographic key pair to an entity by issuing an X.509 certificate (Rescorla and Dierks, 2008), establishing a hierarchical trust model. At the top of the hierarchy is a small number of root CAs whose self-signed certificates (“root keys”) are pre-installed in the trust stores of browsers and operating systems, forming a trust anchor. Any leaf certificate derives its trustworthiness directly or indirectly (via intermediary certificates) from these root keys (Grindal et al., 2025). Hence, the compromise or invalidation of a root key results in loss of trust in all the certificates derived from it (Chuat et al., 2020). Additionally, the browser clients check the X.509 path-validation rules (Boeyen et al., 2008) and revocation status (Koschuch and Wagner, 2015) to ensure a secure connection. Since 2018, major browsers, such as Google Chrome, have additionally required all TLS certificates to be recorded in public append-only Certificate Transparency (CT) logs, enabling the detection of mis-issuance and supporting root store operators in governing the CAs within their trust stores (Laurie et al., 2021).

Root CAs must comply with the CA/B Forum’s baseline requirements’s (BR) (CA/B Forum, 2025) and undergo regular audits before their root certificates are included in trust stores (Grindal et al., 2025). Certificates can be issued at three validation levels, each with distinct requirements for binding endpoints and authenticating entities (CA/B Forum, 2025): (1) Domain Validated (DV) ensures control over the Fully Qualified Domain Name (FQDN) only (e.g., google.com); (2) Organization Validated (OV) adds basic corporate vetting (e.g., Google LLC); (3) Extended Validation (EV) further

involves rigorous legal-entity checks. The non-profit Let’s Encrypt has become the dominant CA for DV issuance via the ACME protocol (Aas et al., 2019), whereas commercial CAs remain the primary issuers of OV and EV certificates (SSLInsights, 2025). For some time period, Browsers had integrated visual indicators for EV certificates, like a green URL bar. However, following (contested) studies according to which the benefits of EV indicators from an end-user perspective compared to mere DV were limited (Felt et al., 2016; Jackson et al., 2007), all major Browsers had removed these indicators by 2020.

2.2. eIDAS Trust Framework

The eIDAS Regulation (EU 910/2014) specifies a set of qualified trust services – *eDelivery*, *eTimestamping*, *eSeal*, *eSignature*, and *website authentication* – to foster “secure electronic interactions” across the DSM (European Commission, 2023b). In contrast to the WebPKI’s industry-led governance model, eIDAS establishes an extended public-law ecosystem. Figure 1 illustrates the layered structure of the WebPKI and the part of the eIDAS ecosystem relevant to QWACs. While both TLS certificates and QWACs are built on the same BR, the eIDAS ecosystem extends the WebPKI by adding regulatory and institutional layers that govern trust services (e.g., QTSPs) and registries (e.g., member state (MS)-trust lists (TLs)). Each MS designates a National Supervisory Body (NSB) maintaining a TL that enumerates accredited QTSPs. The list of trust lists (LoTL)¹ then provides an EU-wide pointer to those MS-NSBs and their TLs.

Driven by concerns over the limited uptake of existing trust services, fragmentation of national electronic identification (eID) schemes, and broader goals of digital sovereignty (Weigl et al., 2022), the EC proposed a far-reaching revision of eIDAS on 3 June 2021 (European Commission, 2021a). The final amendment, Regulation (EU) 2024/1183, was published on 11 April 2024 (European Commission, 2024). The amendment in Article 45 obliges browser vendors to (a) *recognize* qualified certificates for website authentication, and (b) *ensure that the identity data attested in the certificate and additional attested attributes are displayed in a user-friendly manner*. In turn, Annex IV sets the formal requirements for the structure and information content of the corresponding QWACs.

¹See <https://eidas.ec.europa.eu/efda/trust-services/browse/eidas/tls>.

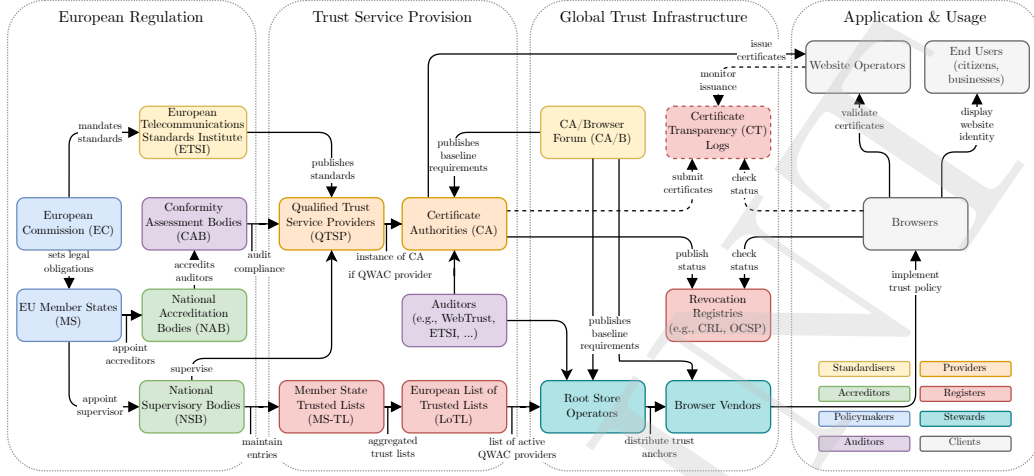


Figure 1: QWAC ecosystem connecting EU regulation, trust service provision, global trust infrastructure, application, and usage.

2.3. Qualified Website Authentication Certificates

According to Article 3(38) of the eIDAS Regulation (European Union, 2014), “*certificate for website authentication means an electronic attestation that makes it possible to authenticate a website and links the website to the natural or legal person to whom the certificate is issued*”. Formally, the regulation is technology-agnostic, yet virtually all current deployments encode the statement as an X.509 object exchanged during the TLS handshake.

QWACs are issued by CAs accredited as QTSPs where each QTSP should satisfy the identity-verification rules in ETSI EN 319 411-1 and undergo a conformity assessment at least every 24 months according to ETSI EN 319 401 and related standards (ETSI, 2024a). This conformity assessment is performed by accredited conformity assessment bodies (CABs), which are in turn assigned by the NSB. QWACs can be deployed and verified in two forms, as defined by ETSI in TS 119 411-5 (v2.1.1) (ETSI, 2025): (a) a 1-QWAC, where the QWAC is used directly as the TLS server certificate (thereby mandating that the issuing CA resides in both Browser trust store and a MS’ TL); and (b) a 2-QWAC binding option, where a conventional WebPKI TLS certificate secures the connection while the website’s identity is bound to a separate qualified certificate issued by a QTSP via the subscriber’s (web server’s) public key. The technical specification defines the required identity data, validation procedures, and – albeit quite vaguely –

browser-facing indicators, establishing a technical baseline (ETSI, 2025).

Cryptographically, both types of QWAC offer the same transport security guarantees as an ordinary TLS certificate. Their distinctiveness lies only in the vetting of organizational identity of the subscriber and the governance framework CAs need to abide by. It is important to note that QWACs policies provide heterogeneous profiles for the subscriber’s identity vetting (ETSI, 2025), ranging from QEVCP (similar to EV TLS certificates) to QNCP-*w-gen*, whose vetting process can be considered weaker than OV (Adriano Santoni, 2025). Moreover, at the time of writing, we are not aware of official implementation guidelines for visual indicators for QWACs, let alone differentiations according to the identity vetting profiles.

2.4. *The QWAC Controversy*

The controversy surrounding QWACs is rooted in their dual traits. Technically, they represent a form of TLS certificates, providing encryption and server-authentication. Their distinctiveness lies in their legal qualification as issuance by a QTSP under EU supervision conveys statutory guarantees of vetting and liability (Entschew and van Brouwershaven, 2024). This legal framing shifts the debate from technical efficacy to questions of governance, accountability, and the extent to which regulatory mandates should override browser-driven trust decisions.

Although QWACs have existed since 2014, their adoption has been mainly limited to financial services and outside Web services, following mandates by the revised PSD2 Directive (EU 2015/2366) (Kudra et al., 2022) and a few eIDAS 1.0 nodes. The scope of browser obligations hence changed significantly with the Commission’s 2021 proposal: Article 45 would require browsers to recognize QWACs and display the attested identity, shifting the debate from optional support to mandated handling (European Commission, 2021a). The controversy intensified in 2023, when a leaked trilogue draft law was interpreted as broadening obligations and narrowing safeguards, prompting coordinated responses from browser vendors and civil society, including an open letter signed by more than 500 scientists and several Non-Governmental Organizations (NGOs) (Mozilla, 2021; Scientists and NGOs, 2023a). This mobilization was led, in particular, by Mozilla under the “*SecurityRiskAhead*” campaign, which published headlines such as “*Last Chance to fix eIDAS: Secret EU law threatens Internet security*” (Mozilla, 2023b). The adoption of Regulation (EU) 2024/1183 then established a legal baseline without technical implementation details, thus leaving room for

uncertainty and speculation. ETSI TS 119 411-5 (v2.1.1), published in early 2025, subsequently provided some clarity by providing technical profiles for consumption models.

The controversy accordingly evolved through successive milestones (proposal, leak-driven escalation, and post-adoption implementation). This temporal context is described in Section Appendix A.1, which provides a timeline of relevant events. The timeline aims to help understand that several prominent arguments were formulated in response to the draft text and corresponding interpretations and ambiguities that differed from the adopted Regulation (EU) 2024/1183.

3. Methodology

This study employs a qualitative mixed-method approach to investigate the technical, human, and institutional dimensions (Orlikowski, 1992) of QWACs within the socio-technical context of eIDAS. Our research design combines a multivocal literature review (Garousi et al., 2016), i.e., a systematic collection and analysis of academic and gray literature, with stakeholder interviews to collect data in the form of (components of) arguments exchanged during the controversy. In doing so, we follow a sequential design: we first assess peer-reviewed evidence through a systematic literature

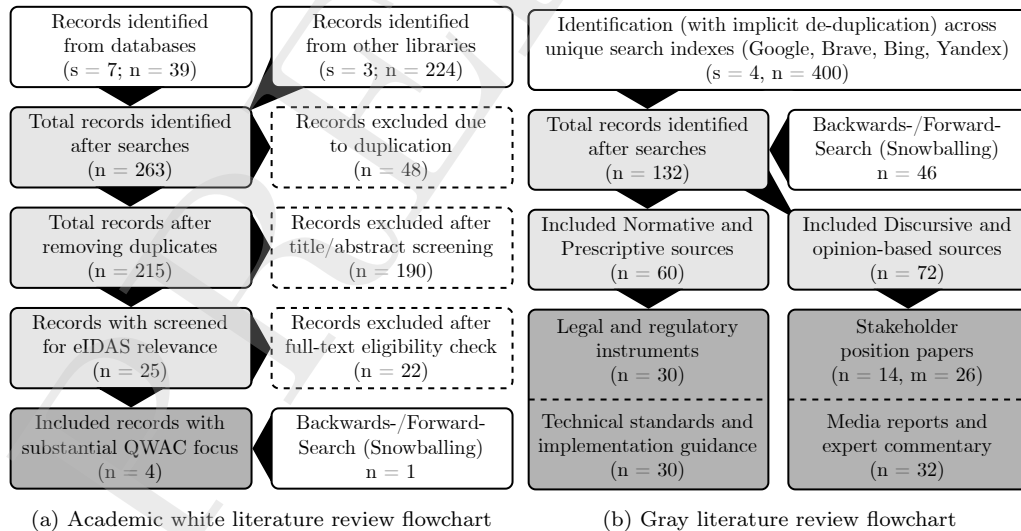


Figure 2: Overview of the multivocal literature review process

review (SLR) (Section 3.1), followed by an extension of the corpus with gray literature to capture stakeholder discourse and documentary evidence (Section 3.2). For all identified records, we conduct a structured analysis of arguments criticizing and supporting QWAC using qualitative coding techniques (Saldaña, 2021) – more specifically, we used a codebook aligned with Toulmin’s model (Toulmin, 2003). We conclude the data collection by conducting interviews (Section 3.4) to triangulate and refine the argument map derived from the multivocal corpus, particularly where documentary evidence is weak or ambiguous. Throughout the analysis, we distinguish a *source-level* orientation (supportive, critical, or neutral reporting) from the stance of an *argument*, which can only be supportive or critical. A “neutral argument” would correspond to descriptive and balanced coverage without explicit advocacy, which can hardly apply for single arguments (or components thereof). “Neutral” is therefore not treated as a third argument stance.

The majority of claims put forward in the QWAC controversy are forward-looking: They assert risks and benefits that depend on future implementation choices, institutional enforcement, and legal interpretation, and therefore cannot yet be verified through outcome-level empirical evaluation alone. For instance, at the time of the open letter (Scientists and NGOs, 2023a), the revision of Article 45 was still in progress, and there was also no implementing act that would detail the technical specifications of QWACs relevant to Browser recognition. We hence use Toulmin’s model as an established argumentation framework to make the underlying assumptions and justifications explicit and to compare stakeholder positions using a consistent structure across sources. This approach ensures a comprehensive exploration of the QWAC controversy, addressing gaps in existing research and narratives as well as providing balanced insights into policy and practice.

3.1. Academic Literature Review

The first stage of data collection was conducted using a SLR (Levy and Ellis, 2006), following the PRISMA guidelines of Moher et al. (2009). We conducted searches in Q4 2024 across 7 academic databases (*ACM Digital Library*, *AIS eLibrary*, *arXiv.org*, *Elsevier ScienceDirect*, *Emerald Publishing*, *IEEE Xplore*, and *Springer*) and 3 additional libraries (*Elicit*, *Dimensions*, and *Google Scholar*) for maximum coverage. The SLR employs a precise search string: "QWAC" OR "Qualified Website Authentication Certificate" OR ("qualified" AND "website" AND "authentication" AND

"certificate" AND "eIDAS"). This process encompassed identification, title/abstract screening, and full-text eligibility assessment. Because our goal was to assess the extent of academic coverage of a niche concept, we intentionally applied a low inclusion threshold: a record passed screening if it was related to eIDAS (as opposed to, e.g., physical quantities Q_{WAC}) and passed eligibility if it contained a substantive discussion of QWACs beyond a passing mention.

As illustrated in Figure 2a, we identified an initial set of 263 records. After de-duplication, 215 unique records remained. During title/abstract screening, 190 records were excluded as not substantively related to eIDAS, leaving 25 records for full-text assessment. Of these, 22 were excluded at eligibility (typically because QWACs were mentioned only in passing), resulting in 3 included academic records. One additional eligible academic record was identified via backward/forward search (snowballing). In sum, our SLR yielded 4 included studies (Entschew et al., 2022; Martius et al., 2024; Entschew and van Brouwershaven, 2024; Wazan et al., 2024). In our mixed-method sequence, the SLR establishes the baseline for coverage of the phenomena in the academic discourse and thereby highlights, notwithstanding the presence of open letters signed by many researchers, evidence gaps that motivate the subsequent systematic inclusion of gray literature.

3.2. *Gray Literature Review*

Due to the limited availability of academic white literature and the extensive coverage of the QWAC controversy in the public discourse and press, gray literature constitutes an essential source of this study. Following the Luxembourg definition (Luzi, 2000), we treat gray literature as documents produced by academia, businesses, governmental bodies, and industry associations that are not commercially published. In our context, this also includes position papers and other documents published by NGOs that have shaped the QWAC controversy.

We conducted keyword searches using web search engines with distinct indexes (Google, Brave, Bing, and Yandex) and screened the top-ranked results to identify entry points. We then expanded coverage through backward/forward snowballing (embedded hyperlinks, back-links, and cross-document references). In addition, we performed targeted searches on LinkedIn Pulse. Across all types of gray literature, we only included content when it constituted a long-form opinion piece or a discussion thread with substantive argumentation. We did not include social media posts or

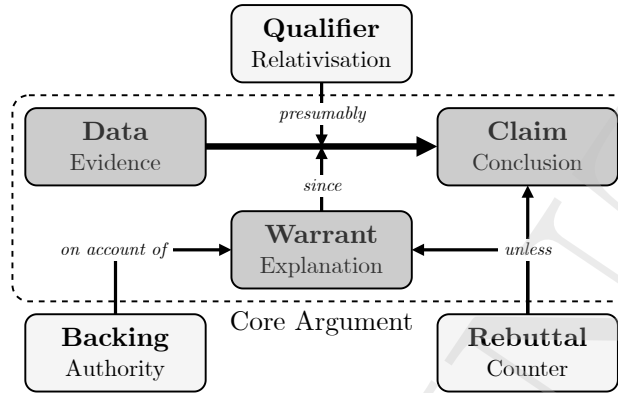


Figure 3: The Toulmin argumentation model adapted from (Toulmin, 2003)

short-form content that does not provide a reasonable amount of relevance according to the AACODS (Authority, Accuracy, Coverage, Objectivity, Date, Significance) checklist (Tyndall, 2010).

The identified gray-literature corpus covers normative, prescriptive documents, including legal and regulatory instruments, technical standards, and implementation guidance (e.g., by ETSI and European Union Agency for Cybersecurity (ENISA)), and root-program and browser-vendor documents (e.g., extracted from the Common CA Database (CCADB) and Mozilla), as well as non-normative, opinion-based documents, such as stakeholder position papers and other web-based sources, including news (journalistic reporting), opinion pieces (blogs, editorials, advocacy texts), and expert discussions (professional networks and fora) (Kuzman et al., 2023). The final set comprises 132 gray-literature records, as summarized in Figure 2b. Together with the white literature, these sources form the multivocal corpus that is coded into Toulmin components to derive the initial argument map (Section 3.3) and to identify publicly engaged stakeholders for interview recruitment (Section 3.4).

Unpacking the QWAC controversy relies on the temporal positioning of arguments relative to the exact wording of the legal and technical specifications available at that time. For instance, during the legislative process, a leaked draft revision of the eIDAS regulation played a key role in the controversy. It is critical to be aware of the evolution of eIDAS and the relationship between Recital 65, Article 45, and Annex IV. Recitals provide the context and reasoning behind the regulation and hence serve as interpretative guides,

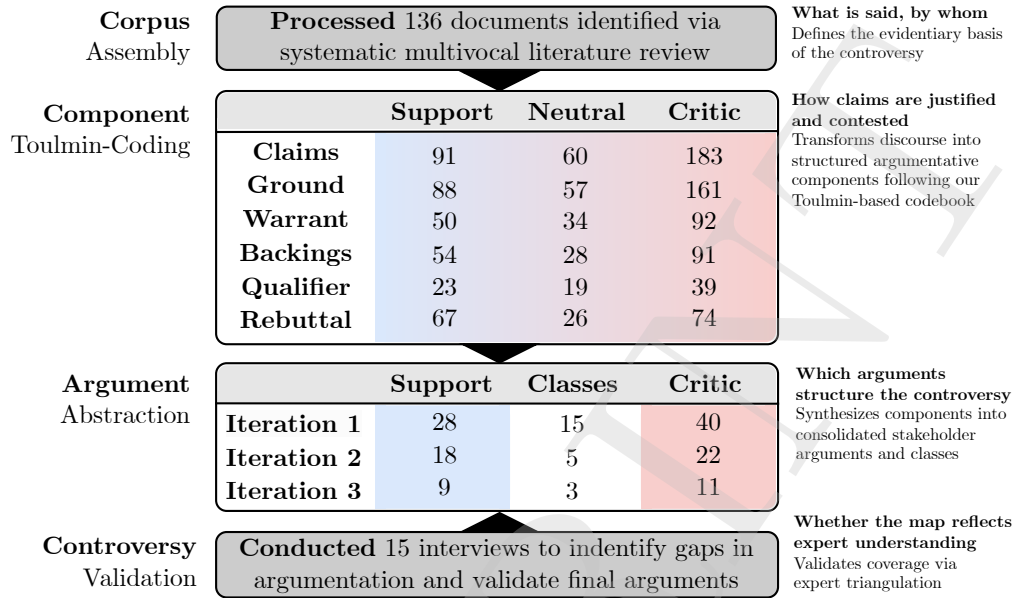


Figure 4: Pipeline to derive the Stakeholder Arguments Map from multivocal discourse

whereas articles constitute the actionable, legally binding parts of the legislation (Klimas and Vaiciukaite, 2008; European Commission, 2015). Because many statements in the corpus respond to different legislative drafts and standardisation milestones, we summarize the key events and successive milestones in Appendix A.1. Unless stated otherwise, we interpret arguments against the final version of the eIDAS Regulation (European Union, 2024) and ETSI technical specification TS 119 411-5 (v2.1.1).

3.3. Argument Abstraction

Our argument abstraction proceeds in three stages, as visualized in Figure 4. Here, the term *argument abstraction* denotes the systematic reduction of many concrete statements into a smaller set of structured arguments.

Stage 1 (*corpus assembly*) comprises the multivocal literature corpus derived from the academic and gray literature reviews described in Sections 3.1 and 3.2. Much of the QWAC controversy consists of forward-looking, contested claims about security, governance, and compliance whose validity depends on future implementation choices, institutional enforcement, and legal interpretation. Therefore, outcome-level evidence was and is largely unavailable. We hence treat argumentation as the primary unit of analysis and use

Toulmin’s model (Toulmin, 2003) to represent each position as a claim with immediate justification and potential rebuttal Liu and Xiong (2024).

Stage 2 (*component coding*) applies a Toulmin-based codebook to decompose corpus statements into argumentative components and to assign stakeholder stances. In this stage, we operationalise Toulmin as a relational and inductive coding framework (Toulmin, 2003; Saldaña, 2021; Liu and Xiong, 2024): extracted text segments are coded as Toulmin components (Figure 3) and linked to capture how stakeholders justify, qualify, and contest claims. For each record, we additionally coded the source stance as supportive, critical, or neutral; neutral sources either present competing positions without endorsing one or report positions without taking a stance. Note that *neutral* is a source-level label (non-advocacy), not a third argument stance; components extracted from neutral sources are attributed to the corresponding supportive or critical arguments in Stage 3. This source-level representation makes implicit technical assumptions, regulatory mandates, and governance expectations explicit, enabling structured comparison across heterogeneous sources without requiring outcome-level evidence (e.g., observed usage or post-deployment attacks).

Finally, Stage 3 (*iterative argument abstraction*) consolidates these coded components into arguments by reducing conceptual overlap and stabilizing a set that comprehensively represents the QWAC controversy. Expert interviews are not considered as a separate stage but are embedded in this final stage as a validation and refinement input to qualify and extend corpus-derived arguments. In this stage, we abstract the coded components into consolidated arguments through three refinement iterations to reduce overlap and stabilize a set with broad discourse coverage. Iterations 1–2 used the literature corpus to progressively consolidate related components into provisional argument classes (40 positions grouped into 5 classes). Iteration 3 integrated the 15 semi-structured interviews by mapping interview statements onto the provisional classes to validate relevance, clarify implicit warrants and qualifiers, and add missing justifications where the documentary record was weak or contested. Throughout, component definitions and emerging argument themes were refined via multi-researcher discussions on consolidation decisions (merging/splitting, naming, and setting boundaries for arguments) to stabilize boundaries with broad coverage and minimal overlap.

The process yielded 20 consolidated arguments emerging from 1237 coded components (334 claims, 306 grounds, 176 warrants, 173 backings, 81 qualifiers, and 167 rebuttals), summarized in Figure 4. The resulting stake-

holder argument map traces where claims are supported, qualified, or rebutted across sources; component counts indicate emphasis within our corpus rather than empirical prevalence. Section 3.4 details the interview protocol and the analysis that provides the validation input used in Iteration 3.

Appendix B is the primary empirical artifact of this study: a Toulmin-structured argument catalogue for all 20 consolidated arguments, with an evaluation template that readers may apply for self-assessment. We do not report author-assigned quantitative strength ratings because they depend on subjective value judgments. In Section 6, we additionally translate the consolidated argument map into an argument-linked threat model that connects stakeholder claims to concrete threats and non-prescriptive design levers.

3.4. Stakeholder Interviews

The last stage of our research design incorporates a set of qualitative, semi-structured interviews (Kallio et al., 2016). This is intended to further extend our evidence base and validate our derived arguments. We identified an initial set of potential participants during the literature review, considering individuals who have been publicly engaged with the QWAC topic, e.g., who signed open letters and have an academic profile in security research, who are involved with governmental bodies, or who are involved in policy-making processes. This group hence comprised high-profile individuals with diverse roles, including regulators, industry experts, and researchers, thus offering a broad spectrum of perspectives. Our attempts to recruit interview partners from browser vendors and the European Commission were unsuccessful, most likely because of the active negotiations in the context of the legislative process. Using purposive expert sampling (Campbell et al. (2020)), we successfully recruited 15 participants (see Figure 5). The study design was reviewed and approved by the Ethics Review Panel (ERP) of our institution, and the interviews were treated confidentially and pseudonymized in compliance with the GDPR (European Union, 2016).

The interviews were conducted predominantly virtually and were subsequently transcribed. On average, interviews lasted 62 minutes, with the shortest lasting 49 minutes and the longest lasting 88 minutes. We used a semi-structured interview guide to ensure consistency while allowing flexibility in addressing emergent themes (Schultze and Avital, 2011; Degen and Teubner, 2024). We thematically analyzed the transcribed responses to identify patterns, contrasts, and novel insights. We used abductive coding (Vila-Henninger et al., 2024) to refine the codes obtained from the initial induc-

tive coding of the academic and gray literature. These findings enrich our argumentation analysis by highlighting aspects that have not been publicly discussed previously. We employ two code systems to capture both the interview structure and the argumentative coverage. First, we applied a topic code system based on the semi-structured interview guide (e.g., problem framing, perceived risks, governance capacity, deployment pathways) (see Figure A.2), which allowed us to trace how often and in what depth each interviewee addressed the key questions. Second, we applied the stage 3 arguments as a code system to validate coverage. The analysis addressed several major and minor themes, the frequencies of which are visualized in Figure A.3, which also serves as an outline of topics and questions asked during the interviews. For the argument code system the figure depicts how often each of the 20 arguments appeared in the interview corpus. While illustrating the concentration of discussion points across interviews, the participants are grouped by their stance, which helped us identify which arguments are primarily driven by supportive policy actors, which ones by browser-aligned critics, and which ones are shared across groups.

Interviews were conducted specifically to triangulate how stakeholders interpret different aspects concerning QWACs, what arguments they bring forward, and to refine and extend Toulmin components (e.g., grounds, warrants, qualifiers) where the documentary record was weak or ambiguous. We then mapped interview insights back onto the argument map to indicate whether they confirm, contradict, or qualify the corpus-derived evaluation.

4. Analytical Framework

In this section, we present a preliminary descriptive analysis of the controversy based on the multivocal literature review and interview data. Based on the axial coding of 20 core arguments, as depicted in Figure 6, we aggregated the argumentative material into three heuristic and partially overlapping categories: *Security & Trust (ST)*, *Governance & Authority (GA)*, and *Compliance & Interoperability (CI)*, describing recurrent arguments using Toulmin’s model.

4.1. Stakeholder Typology

We classify stakeholders into three categories (supporter, critic, neutral) based on their positions throughout the analysis (see Figure 5): (a) supporters argue in favor of QWACs or the amended Article 45; (b) critics argue

ID	Stance	Domain	Experience	Engagement
E01	Supporter	Industry	18 years	High
E02	Neutral	Academia	30 years	Low
E03	Critic	Academia	8 years	Moderate
E04	Critic	Academia	30 years	High
E05	Neutral	Governmental Agency	30 years	Moderate
E06	Critic	Academia	15 years	Low
E07	Critic	Academia	35 years	High
E08	Supporter	Regulatory Body	13 years	High
E09	Supporter	Service Provider	25 years	High
E10	Supporter	Service Provider	20 years	High
E11	Neutral	Service Provider	25 years	High
E12	Neutral	Service Provider	33 years	High
E13	Critic	Academia	8 years	Low
E14	Supporter	Governmental Agency	42 years	High
E15	Supporter	Standards Organization	26 years	High

Figure 5: Interviewee engagement in the discourse by stance and domain (random order)

against it; and (c) neutrals present observations or comparisons rather than explicit advocacy. Stances are inferred from each stakeholder’s authored or endorsed source materials and, where applicable, interview statements. The label captures an overall orientation for aggregation, not uniformity across communications, and actors may endorse isolated arguments while overall favoring one stance. We apply the neutral stance only at the source level, not argument-level. Arguments always have a stance (i.e., supporting or criticizing), whereas a truly neutral argument can be understood as one entirely free from values, perspective, and framing, which is not suitable for this study.

Section 3.3 reports the consolidated argument set and coded-component totals. Here, we focus on stakeholder affiliation and visibility to contextualize which actor types most prominently advance or contest the main positions. Three iterations of classification refinement led to 9 core arguments for and 11 arguments against QWACs, as presented in Figure 6.

The stakeholder analysis revealed a diverse set of actors. Supporters included influential organizations such as the European Signature Dialog (ESD) (an NGO representing QTSPs), prominent CAs, ETSI, and Bundesdruckerei Gruppe GmbH – the German Federal Printer, all of whom played a key role in shaping the narrative in favor of QWACs. The ESD, in particular, strongly supported the mandatory browser trust for QWACs, emphasizing its poten-

tial to bolster EU digital sovereignty and enhance consumer trust. European institutions (e.g., the EC) also played a central role in the pro-QWAC argument, advocating for revisions to Article 45 as a necessary step toward regulatory alignment and enhanced web security in the EU. Moreover, CAs, such as D-Trust (which belongs to the German Federal Printer) and Entrust, voiced for regulatory harmonization, emphasizing interoperability, usability, and trust in digital identity systems, as well as the lack of lawfulness in browsers’ root store decision-making. While both D-Trust are QTSPs entitled to issue QWACs, only D-Trust has roots included in major trust stores, as Entrust roots were removed from most trust stores in late 2024.

In contrast, Mozilla has been the most visible critic of QWACs, warning of the privacy risks inherent in a state-controlled certificate system and even staging a public protest in Brussels (Mozilla, 2023a). The Electronic Frontier Foundation (EFF), along with scientists and NGOs, echoed these concerns, highlighting the risks of government surveillance through the interception of encrypted traffic and the loss of independent oversight. Their position was formalized in an open letter signed by more than 500 scientists worldwide (Scientists and NGOs, 2023a).

4.2. Analytical Themes

Drawing on the socio-technical theory of the duality of technology (Orlikowski, 1992), we treat QWACs as a socio-technical intervention whose implications are enacted by stakeholders under institutional rules (eIDAS) and infrastructure governance (WebPKI). Rather than separating technical design from social impact, as a result of the argument abstraction and analysis (Section 3.3), we organize the controversy into three themes – *Security & Trust*, *Governance & Authority*, and *Compliance & Interoperability*. This reflects how stakeholders argue from a disputed mechanism to an invoked consequence. In each theme, the first term represents the dominant technical mechanism under legal mandate (security, governance, compliance), while the second term infers the socio-institutional stake that the argument claims this mechanism affects (trust, authority, interoperability). The themes are used as a means to better organize, present, and compare the final set of arguments, but do not imply that arguments are confined to a single theme. The Toulmin decomposition for each consolidated argument is provided in Appendix B.

ST *Security and Trust*. Supporters argue that QWACs strengthens website identity assurance by binding TLS protection to legally accountable

identities (ETSI, 2025). They emphasize that qualified status and conformity audits provide standardized and transparent trust indicators in browsers, designed to make organizational identity visible and auditable (ESD, 2023c; Entrust Corporation, 2021) that can help mitigate phishing and fraud (Drury and Meyer, 2019).

Critics counter that QWACs do not meaningfully enhance user risk awareness, since prior studies show that identity indicators fail to change behavior (Jackson et al., 2007; Thompson et al., 2019). They add that mandatory trust signals may not deliver effective guarantees of trustworthiness, as legal assurances cannot compensate for technical flaws in issuance (Mozilla, 2020; Electronic Frontier Foundation (EFF), 2022). Concerns extend to degraded data protection (Claburn, 2023; Mozilla, 2023a) and the introduction of new attack vectors if compromised government-endorsed CAs are included (Keizer, 2011; Raman et al., 2020).

GA *Governance and Authority*. Supporters assert that QWACs promote fair competition in digital markets by curbing the dominance of foreign browser vendors in setting trust policies (Bitkom, 2022; ESD, 2022b). They argue that embedding trust in MS-TLs strengthens EU digital sovereignty, shifting governance from private corporations to statutory supervision (Entschew and van Brouwershaven, 2024; Entschew et al., 2022). In this perspective, QWACs also advance the EU DSM by extending the European trust space across borders and harmonizing identity assurance (European Parliament, 2024).

Critics respond that Article 45 could facilitate government surveillance through the state-controlled certificate issuance (Mozilla, 2023b; Electronic Frontier Foundation, 2023). They argue that mandatory recognition undermines the global trust models coordinated by the CA/B (Grindal et al., 2025; Braun et al., 2014). Moreover, critics highlight that statutory requirements may conflict with existing root store standards and practices, potentially constraining browser mandates for CT, incident response, and revocation policies (Helme, 2023; Mozilla et al., 2023; Rescorla, 2022).

CI *Compliance and Interoperability*.

Supporters stress that QWACs integrate website authentication into

the broader EU trust framework, aligning it with other qualified services (ETSI, 2023c; European Commission, 2021a). They argue that this integration supports cybersecurity directives (e.g., the Revised Network and Information Systems (NIS2) directive) and harmonized supervision (Wazan et al., 2024). By binding legal identity to web services, QWACs are said to strengthen accountability and transparency across financial services and e-government contexts (LuxTrust S.A., 2024).

Critics counter that mandatory QWAC recognition risks fragmenting the global WebPKI by creating parallel trust ecosystems with divergent handling of validity and revocation (CCADB, 2020a; Scheitle et al., 2018). They caution that enforcing EU-specific standards undermines the technological neutrality principles central to the evolution of the WebPKI (Mozilla, 2021; Delignat-Lavaud et al., 2014). Others point to the increased complexity and operational costs for website operators, who must manage multiple profiles and audit regimes (Tehrani et al., 2024; Martius et al., 2024), particularly when considering the many connections maintained when interacting with a modern web service (CCADB, 2020a). Finally, critics stress that QWACs underperform compared to existing measures such as automated monitoring, which already improves accountability without statutory compulsion (Laurie, 2014; Certificate Transparency, 2020).

4.3. Shared Premises

To ensure consistency in our evaluation, we distinguished between two layers of components. First, *shared components* are field-invariant: they define the stable reference frame within which all claims are assessed, and second, *argument-level components* are field-dependent: they vary with the substance of each claim and specify what counts as acceptable grounds, warrants, or backing within the argument’s context (Toulmin, 2003). This means that all arguments, regardless of class or stance, operate within this structural frame.

Identifying the shared components explicitly allows us to avoid redundancy and ensures that all claims are tested against the same underlying criteria. This makes the comparative evaluation consistent because the basis for the judgment is declared in advance rather than inferred for each argument. The list below outlines the shared premises (referenced as Px), whereas the field-dependent components are detailed in Appendix B.

- P1** The *Legal Foundation* defines the binding legal framework provided by eIDAS 2.0 (EU/2024/1183), primarily Articles 45 and Annex IV.

All claims are interpreted against the final legislative text, rather than (leaked or published) regulation drafts.

- P2** The *Technical Specifications* define the operative certificate and consumption models as codified in ETSI TS 119 411-5 (ETSI, 2025) and EN 319 401/412 (ETSI, 2024a). This implies that all evaluations of feasibility, interoperability, or attack surface are judged relative to the defined 1-QWAC and 2-QWAC profiles (ETSI, 2025).
- P3** The *Institutional Governance* defines the EU trust-service framework, including MS-TL/LoTL, NSB, and CAB-related processes. This means that claims about recognition, supervision, or conformity are assessed in terms of how authority is exercised within the eIDAS institutional framework (Entschew and van Brouwershaven, 2024; Martius et al., 2024; Wazan et al., 2024).
- P4** The *Operational Enforcement* defines how recognition and distrust decisions are executed through the BR governed by the CA/B, root-store policies, including CA admission, compliance monitoring, and incident handling. This means that claims are assessed in terms of how these root programs implement and enforce trust anchors in practice (CA/B Forum, 2025; Mozilla, 2023c; CCADB, 2018, 2020a).

5. Argument Map

This section summarizes the consolidated arguments by classes. The arguments along the previously introduced classes are outlined in Figure 6, for which the primary empirical artifact is provided in Appendix B.

5.1. Security and Trust

- ST-S1** QWACs *strengthen Website Identity Assurance* by binding a supervised legal-entity identity to a TLS endpoint through harmonized semantics in EN 319 412 and the consumption models in ETSI TS 119 411-5 (ETSI, 2025, 2023b, 2024a). This uplift is governance- and process-based rather than cryptographic, but it is consistently supported across regulation, standards, and supervisory practice (ENISA, 2016).
- ST-S2** QWACs *display a Transparent Trust Indicator* by mandating that browsers provide verified identity fields in a user-friendly manner, as required by Article 45 and ETSI EN 319 412-4 (ETSI, 2025). Harmonized

semantics exist, yet empirical studies are missing that confirm advantages over deprecated EV indicators (Felt et al., 2016).

ST-S3 QWACs *protect Users from Fraudulent Websites* by exposing vetted legal names and jurisdictions that can be compared to brand claims at connection time (ETSI, 2025; ENISA, 2017). Standards ensure attributes are retrievable, yet phishing research shows users often ignore or misinterpret identity signals (Jackson et al., 2007; Drury and Meyer, 2019).

ST-C1 QWACs *do Not Enhance User Risk Awareness* because most hu-

Security & Trust

ST-S1	Strengthen Website Identity Assurance
ST-S2	Display Transparent Trust Indicator
ST-S3	Protect Users from Fraudulent Websites
ST-C1	Do Not Enhance User Risk Awareness
ST-C2	Do Not Deliver Effective Trust Signals
ST-C3	Degrade User Data Protection
ST-C4	Introduce New Attack Vectors

Governance & Authority

GA-S1	Promote Fair Competition in Digital Markets
GA-S2	Strengthen EU Digital Sovereignty
GA-S3	Advance the EU Digital Single Market
GA-C1	Facilitate Government Surveillance
GA-C2	Undermine Neutral Global Trust Models
GA-C3	Conflict with Existing Root Store Standards

Compliance & Interoperability

CI-S1	Integrate Website AuthN into EU Trust Schemes
CI-S2	Integrate with EU Cybersecurity Directives
CI-S3	Strengthen Accountability and Transparency
CI-C1	Create Fragmented Trust Ecosystems
CI-C2	Undermine Technological Neutrality Principles
CI-C3	Increase Complexity and Costs for Website Operators
CI-C3	Underperform Compared to Existing Measures

Figure 6: QWAC arguments grouped by class and stance. Argument IDs follow *CLASS-STANCE+N* (e.g., *ST-S3*). Stance colors: blue = Support (S), red = Critic (C).

mans do not adapt their behavior when shown certificate cues, and removing EV indicators did not reduce overall security (Felt et al., 2016; Hunt, 2019; Mozilla, 2020). Given the aged, limited evidence on EV, based on existing studies of user behavior, harmonized QWACs could still prove useful for automation and training.

ST-C2 QWACs *do Not Deliver Effective Trust Signals* because certificate identity cues themselves are too weak to influence reliable decisions by untrained users (Biddle et al., 2009; Rescorla, 2022). Some evidence suggests that standardizing the visualization of identity cues, thereby mitigating their current ineffectiveness, may lead to little learning (Felt et al., 2016; Thompson et al., 2019).

ST-C3 QWACs *degrade User Data Protection* by introducing additional on-line discovery or revocation/status checks that may expose user browsing metadata to QTSPs or related registries, especially in the 2-QWAC binding model (CCADB, 2020a; Mozilla, 2022). However, neither the regulation nor the standards mandates per-visit *phone-home* validation, and stapled or otherwise offline status mechanisms can avoid such metadata leakage (Berbecaru and Liroy, 2023).

ST-C4 QWACs *introduce New Attack Vectors* by introducing additional verification logic and mandated UI elements, increasing complexity and the risk of user over-trust (Rescorla, 2022; Felt et al., 2016). These risks remain theoretical, with no incident evidence to date mitigating their weight (CA/B Forum, 2025; Helme, 2023).

5.2. Governance and Authority

GA-S1 QWACs *promote Fair Competition in Digital Markets* by creating a statutory path for QTSPs to be recognized in browsers via MS-TL and LoTL listings (Bitkom, 2022). Position papers argue strongly for competition, but empirical evidence of changed market shares is absent (Bundesdruckerei, 2022; Bundeskartellamt, 2022), and neither form of QWACs as ultimately specified by ETSI TS 119 411-5 supports this proposition.

GA-S2 QWACs *strengthen EU Digital Sovereignty* by embedding organizational authentication into EU-supervised trust frameworks (European Commission, 2021c; Entschew and van Brouwershaven, 2024). The Regulation and ETSI baselines explicitly shift governance to statutory supervision, even if browsers retain technical power (Mozilla, 2023c).

- GA-S3** QWACs *advance the EU Digital Single Market* by harmonizing recognition and display of identity across borders (Weigl et al., 2022; European Parliament, 2024). Standards provide mechanisms, but measurable DSM efficiency gains remain unproven (ETSI, 2023d,b).
- GA-C1** QWACs *facilitate Government Surveillance* because state-controlled QTSP could issue interception QWACs that must be accepted until detected (Mozilla, 2023a; Raman et al., 2020; Amann et al., 2017). The mechanism is plausible and historically evidenced, but mitigations like CT logging mitigate the risk by making such attempts publicly observable (Laurie et al., 2021; Claburn, 2023). While stakeholders broadly expressed their support for mandatory CT logging in the interviews, the formulation in Article 45 could be understood in a way that would restrict CT logging (Helme, 2023; Scientists and NGOs, 2023a), and to date, the technical specification of QWACs remains silent on this issue.
- GA-C2** QWACs *undermine Neutral Global Trust Models* because mandatory legal recognition constrains browser discretion and may conflict with CA/B Forum norms (CCADB, 2020b; CA/B Forum, 2020; Grindal et al., 2025). Alignment work exists, but friction remains foreseeable (ETSI, 2025).
- GA-C3** QWACs *conflict with Existing Root Store Standards* by shifting trust anchor criteria from browser policy to statutory listings (Mozilla, 2021; Internet Society et al., 2023). Politicized inclusion remains a risk primarily of the 1-QWAC approach, whereas the 2-QWAC model would continue to rely primarily on existing trust roots (Entschew and van Brouwershaven, 2024).

5.3. Compliance and Interoperability

- CI-S1** QWACs *integrate Website Authentication into EU Trust Schemes* by standardizing identity validation under ETSI supervision and MS-TL/LoTL governance (ETSI, 2024a; Wazan et al., 2024). The standards pathway is coherent, but deployment evidence is limited.
- CI-S2** QWACs *integrate with EU Cybersecurity Directives* by aligning with ENISA recommendations and NIS2-related governance (ENISA, 2017; European Commission, 2020b). This situates QWACs within a coherent policy framework, although the distinct benefits of such integration beyond existing standards remain debated (European Commission, 2021b; Mozilla et al., 2023).

- CI-S3** QWACs *strengthen Accountability and Transparency* by surfacing legally verified identities at the point of interaction, consistent with GDPR principles (Bundesdruckerei, 2022; Bailey, 2022). Accountability is mandated in law, yet transparency at the interface remains questionable (Martius et al., 2024; Bailey, 2022).
- CI-C1** QWACs *create Fragmented Trust Ecosystems* because 2-QWAC bindings require extra discovery and status checks not covered by existing automation (CCADB, 2020a; Mozilla, 2021). Standards aim for alignment, but vendor adoption patterns remain uncertain.
- CI-C2** QWACs *undermine Technological Neutrality Principles* by privileging one artifact and narrowing space for alternative solutions (Internet Society et al., 2023; Mozilla et al., 2023). The Regulation is formally agnostic, but critical analyses foresee constraints on innovation (ETSI, 2025).
- CI-C3** QWACs *increase Complexity and Costs for Website Operators* by requiring dual audits and new verification tooling for operators (CCADB, 2020a; Rescorla, 2022). An expert commentary notes potentially high fixed costs, though economies of scale may mitigate them (Helme, 2021).
- CI-C4** QWACs *underperform Compared to Existing Measures* because DV, CT logging, and short lifetimes already provide efficient safeguards (Laurie, 2014; Aas et al., 2019). Critics see QWACs as added complexity, absent outcome gains, while supporters stress the value of legal accountability (Rescorla, 2022; ETSI, 2025).

6. Synthesis and Discussion

In this Section we synthesize the argument map (Figure 6) with interview triangulation (Section 6.2) to derive theoretical and practical implications (Section 6.3). Appendix B provides the complete argument catalogue and evaluation framework.

6.1. Synthesis

Most claims across the three argument classes were logically sound but conditional. Supporters ground their case in the binding of legal-entity identity required under eIDAS 2.0 and ETSI TS 119 411-5, while critics emphasize the history of browser trust store governance, the absence of novel security mechanisms, and a lack of evidence that visual indicators have a positive

security impact on end-user behavior (ETSI, 2025; Felt et al., 2016; Thompson et al., 2019).

Positive outcomes appear more plausible when deployments favor the 2-QWAC model to preserve existing WebPKI chains, mandate CT logging for qualified artifacts, adopt privacy-preserving revocation mechanisms such as stapling or CRLite, and provide consistent, machine-readable identity semantics through vendor User Experience (UX) hooks (ETSI, 2025; Laurie et al., 2021; Berbecaru and Liroy, 2023; CCADB, 2020a). Negative outcomes become likely if 1-QWAC certificates displace current chains or if supervisory and browser processes diverge, creating fragmentation, governance frictions, and potential for governmental interception of web traffic without demonstrable security gains (CCADB, 2020b; Mozilla, 2021; Grindal et al., 2025). While a mandate to recognize QTSP-issued QWACs as an alternative form of TLS certificates with a mandate for Browser-side recognition was initially marketed to provide a strong form of “European digital sovereignty”, the result of the negotiations led to 1-QWACs requiring also inclusion in existing trust stores, which leads to limited benefits while causing major discomfort among Browsers and security researchers. As such, considering the political capital invested in the controversy as sunk costs and the technical specification of QWACs as given, as well as ongoing discomfort expressed by Browsers and ambiguity as to the implementation of “user-friendly visual indicators”, the friction to implement 2-QWACs may be lower while providing most of the anticipated benefits, contingent on the decisions made.

Because rigorous field evidence remains scarce, effectiveness should ultimately be assessed through pre-registered studies using metrics such as phishing click-through and credential submission rates, mis-issuance detection latency, distrust times via CT, and cross-border onboarding performance. Until such data exist, QWACs are best understood as a governance intervention whose value depends less on cryptographic innovation than on supervision quality and cross-ecosystem coordination (Laurie, 2014; Scheitle et al., 2018; Laurie et al., 2021).

6.2. Interview Triangulation

We use interviews to triangulate and validate the argument map as they clarify stakeholder intent, make implicit qualifiers explicit, and highlight where the documentary record is weak or ambiguous. We added further insights in Appendix A.2. The code matrix analysis of the 15 expert interviews in Figure A.3 reveals a strongly structured debate around QWAC,

with clear differences in argumentative emphasis across supporters, critics, and neutral stakeholders. Critics predominantly mobilized security and governance risk frames, most notably, that QWAC do not deliver effective user trust signals and may facilitate government surveillance, both of which appear with high coding density across multiple critic interviews, indicating sustained and shared concern rather than isolated objections. In contrast, supporters consistently emphasized benefits related to strengthening website identity assurance and protecting users from fraudulent websites, with these codes appearing frequently and across nearly all supporter interviews, suggesting a coherent pro-QWAC narrative focused on institutional trust and assurance. Neutral experts occupied an intermediate position, engaging with both benefit and risk-related codes but with lower overall intensity, often framing concerns in terms of implementation complexity and integration with existing cybersecurity frameworks rather than principled opposition.

The interviews converged on four practical questions: *what problem QWACs are meant to solve, what risks they could create, what governance capacity is needed to balance both, and how deployment could proceed without overburdening relying parties (RPs)*. In the following, we synthesize the most actionable insights from the corresponding discussions, illustrating them with evidence from interview participants.

Experts largely view QWACs as a means to provide a reliable, machine-readable organizational identity rather than a visual cue for everyday users. The primary audience, they argued, is trained operators who can interpret harmonized identity attributes and integrate them into audits or automated security workflows. Participants emphasized that identity data must be measurable and consistently presented across vendors to be valuable, cautioning that past attempts at eye-catching indicators for EV TLS certificates misled users and should not be repeated (E9, E10, and E12). E10 stated: *“I would have considered it more sensible to continue to visualize this organizational affiliation, but in clearer language and in such a way that [...] it is made clear to them that transport encryption exists and that you are communicating with the next party.”* Building on this perspective, interviewees proposed a practical rollout to deploy QWACs, facilitating accountable, machine-readable identity verification where it is currently missing, maintaining existing TLS transport while bindings and status mechanisms mature, and reinforcing the system with CT logging and a robust escalation channel. Success, they advised, should be gauged with operator-centric metrics – not end-user interface sentiment – while debates over EU sovereignty and QTSPs-only recognition

continue among stakeholders willing to accept higher integration costs (E5, E8, E11–E13, and E15). In this context, interviewees also highlighted situations where QWACs could add value outside the traditional browser bar. As E11 noted, “[a] major potential issue is that when I photograph a QR code, I do not really know at that point where the data is going, and there is a link in there that gets evaluated. If I were to secure it with a QWAC and the software displayed it as trustworthy, similar to a browser, then we would also gain security here.” QR codes exemplify low-visibility contexts where surfacing qualified identity information could make a difference, as users today have limited visibility into where the link directs their data to.

A related concern that was frequently raised by critics of QWACs points to frictions with the design of modern web apps, which often rely on dozens of background calls to third-party servers, raising the question whether only the origin or all services need to carry a QWAC (CCADB, 2020a) to make a visual indicator applicable. In such cases, a QWAC signal could indeed introduce major complexities and costs or provide a misleading sense of assurance if the supporting infrastructure is opaque or uncontrolled. Several interviewees, therefore, suggested that any deployment must clarify how qualified identity applies not just to the entry point but to the broader service composition. E8 compared this to car manufacturing by stating “*It’s that when I buy a car, I don’t care who manufactured the wheel, the tyre and the windshield, there is one entity which is responsible for the parts that they are providing to me. Therefore, it is their responsibility to ensure that everything is secure, and if I have an issue, I go to them. I don’t need to identify this sub-service or this component.*” Accordingly, even when acknowledging that modern web services integrate connections to third parties CCADB (2020a), the ultimate responsibility for safety and compliance rests with the provider of the origin, so the supply of a QWAC for the origin should be sufficient. Although three (E9, E10, and E11) out of five experts supported this argument, E7 considered it a very technical matter that required further investigation with the vendors. Interestingly, E15 questioned this analogy for a critical case of an attack, “*I think the difficulty comes into play when it’s not like someone could plant a component within the car without the manufacturer knowing it. And then the question is who’s responsible, like your car is parked and I put something in your car and that breaks it. And that’s where, if anyone within the ecosystem were compromised. Like you trust all or nothing.*”

Multiple experts have described a real gap between statutory recognition and existing root-store governance. They warned that if recognition is read

as indirect pressure on browser inclusion or if Member State listings become de facto trust anchors for identity signaling, incident response could be politicized or slowed (E2, E3, and E6). The same group argued that the legal duty to “*recognize and display*” only yields legitimacy if coupled with operational power. One concrete recommendation was to establish a single point of contact with a European authority (e.g., ENISA) to coordinate precautionary measures across borders and vendors so that mis-issuance or abuse triggers the same fast path regardless of which QTSP is involved (E11). Practitioners were also frank about their capacity and lead times. Audit and integration steps sit on top of what large enterprises already operate, so even supportive teams forecast multi-quarter programs before benefits materialize (E15). Smaller providers and QTSPs face entry barriers from additional audits to be allowed to issue QWACs, automation work for discovery and status, and the need to support RP tooling that does not yet exist off the shelf (E5). These notes explain why some interviewees prefer minimal-friction consumption paths that reuse existing WebPKI plumbing while still surfacing qualified identities (E13).

However, any design choice that lowers integration cost can also weaken the political message some stakeholders want, as one of the key motivations for the development of QWACs in eIDAS 1 originated from the DigiNotar incident, where the Dutch government relied on the goodwill of Microsoft to delay distrust to maintain the availability of key government services (van der Meulen) (E14). E1 supported the claim as follows: “*Diginotar was the reason eIDAS was implemented... Supervisory and conformity assessment bodies now ensure such incidents don’t happen.*” One expert argued that side binding, as implemented in the 2-QWAC approach, keeps browsers in control and therefore does little for a sovereignty narrative; another countered that pushing identity through existing rails is precisely what makes deployment feasible in the near term (E8, E13). This disagreement is less about technology than about which incentives the scheme should prioritize.

Despite differing views on broader policy questions, participants converged on two safeguards as essential for trustworthy QWACs deployment. First, issuance transparency, where several experts advocated mandatory CT logging to deter mis-issuance and strengthen public oversight, with E7 noting that, “*I think Europe could actually make CT mandatory*”. Others, however, cautioned that any CT requirement must remain compatible with the EU principle of technological neutrality (E5). Second, for privacy-preserving validation, interviewees stressed that status checks should rely on stapled

or otherwise offline mechanisms, avoiding per-visit lookups that could leak user-browsing metadata, particularly when QWACs are consumed in mobile or embedded WebView contexts (E11 and E12).

6.3. Theoretical and Practical Implications

The synthesis of the three themes described in our results as Security & Trust, Governance & Authority, and Compliance & Interoperability, together with the identified stakeholder stances and the consolidated argument map, builds a foundation for the broader theoretical and actionable implications. We derive implications by contrasting supportive and critical arguments within each theme and by integrating interview insights that confirm, contradict, or qualify the arguments derived from the literature corpus.

6.3.1. Theoretical Implications

The QWAC controversy is primarily a governance-of-security conflict rather than a cryptographic dispute. Supporters and critics highlight concerns on authority, recognition obligations, and supervisory control ([GA-S2], [GA-C2]), but barely point out TLS-related cryptographic properties. This finding reinforces a socio-technical perspective, highlighting that security outcomes depend as much on institutional arrangements as on technical designs. This implication is also backed by our interviewees, as such E8, who noted that “*QWAC is not special at the technical level, it’s just a normal TLS certificate issued by a CA very normally. The main differences are that QWAC adds some more security, which is not again technical, but which is more at the organizational level.*”

Visual display of QWACs may not enhance user-facing security. The claims about user-facing security benefits ([ST-S2], [ST-S3]) lack empirical evidence. This demonstrates that security debates in regulatory transitions often rely on incomplete evidence, where governance mechanisms can be well-defined in legislation or standards, whereas their actual security effects remain speculative. The contested research on EV indicators even suggests that there may be a chicken-and-egg-problem, where security benefits only have a chance to emerge once they are broadly and consistently adopted, in combination with digital literacy efforts. This insight extends theoretical discussions on the limitations of evidence-based security policymaking.

Legal mandates act as security backings and shift what counts as evidence in security governance. In the QWAC discourse, supportive positions are

frequently grounded in statutory supervision, liability, and harmonized compliance mechanisms, while critical positions emphasize empirically uncertain behavioral effects and global interoperability constraints. This suggests that when law is deployed as a security mechanism, the controversy pivots from protocol-level properties to institutional legitimacy, operational enforceability, and the evidentiary gap between policy intent and measured security outcomes.

6.3.2. Practical Implications

Identity consumption matters more than identity minting. Arguments highlight that QWACs challenges pertain less to how identity is created and more to how it is consumed, interpreted, and displayed in end-user contexts. Supportive arguments such as [ST-S1], [ST-S2], and [CI-S3] emphasise the added value of legally supervised identity and harmonised identity semantics for consumption and display, whereas critics stress that identity cues can be ignored or misinterpreted in practice [ST-C1], [ST-C2]. A distinct critical concern is that some consumption patterns – especially 2-QWAC variants requiring additional discovery or status lookups – could introduce browsing-metadata leakage if not implemented with privacy-preserving, offline validation [ST-C3]. The 2024 settlement and the latest ETSI guidance enable the use of QWAC identity data without altering browser security decisions. This shifts practical focus toward establishing a minimal, stable, and machine-readable set of qualified attributes, ensuring cross-browser and cross-application consistency in identity display, and supporting automated extraction and interpretation of identity fields based on globally recognised organizational identifiers.

Security and trust arise from operational controls rather than labels alone. Arguments in favour of QWACs adoption emphasise accountability and formal supervision [GA-S1, CI-S3], whereas critics stress that user-visible identity labels alone do not provide meaningful security guarantees [ST-C1], [ST-C2] and that additional verification logic can expand attack surface and the risk of over-trust [ST-C4]. Across discussions, stakeholders recognise that the backbone of WebPKI trust today consists of CTs logging, timely and privacy-preserving revocation, and clear incident-response workflows. Therefore, QWACs must integrate with CT – or with European logging services if sovereignty is taken seriously (Fiedler and Thiel, 2014) – and revocation handling must be both robust and privacy-preserving, avoiding per-visit status lookups that could leak browsing metadata [ST-C3]. Moreover, opera-

tional policies should prioritise rapid detection, disclosure, and mitigation of mis-issuance events. This shifts emphasis from symbolic identity signals to verifiable operational trustworthiness.

Institutional roles of regulators and browsers must remain clearly separated. Arguments around governance and authority repeatedly emphasize the importance of maintaining a separation between legal identity qualification and technical governance of trust anchors [GA-C2, GA-C3, CI-C1]. This separation concern is repeatedly framed as a precondition for avoiding parallel trust paths and inconsistent validation and incident-response behavior across user agents. The 2024 compromise avoids imposing new trust anchors and allows browsers to continue applying their own security controls. Practically, this implies that qualified identity information may be displayed, without influencing connection security or path-validation decisions. Additionally, future regulatory updates must preserve this separation to avoid fragmentation and maintain global interoperability.

6.3.3. Actionable Recommendations

Based on the consolidated argument map, interviews, and the threat-model levers (Table C.1), we recommend the following minimum actions per stakeholder group. *Browsers and other user agents* should expose a minimal, interoperable qualified-identity interface (UI and/or API) that is explicitly decoupled from transport-security indicators and validation decisions, and evaluate effects with privacy-preserving telemetry and controlled tests. *Providers* (QTSPs/CAs and status infrastructure) should ensure issuance transparency (e.g., CT logging where feasible), implement privacy-preserving revocation/status (prefer stapled/offline mechanisms over per-visit lookups), and maintain rapid incident-response procedures. *Regulators and standardisers* should preserve the separation between legal qualification and root-store governance and fund reference tooling, conformance suites, and cross-border escalation coordination. *Relying parties* should pilot QWACs consumption in regulated contexts (e.g., revised Payment Services Directive (PSD2) services (EU 2015/2366)), integrate qualified attributes into audit and fraud detection workflows, and report measured outcomes and integration costs back into standards and governance.

6.3.4. Argument-linked threat model

To connect the controversy to implementation decisions, we derive an argument-linked threat model from the consolidated argument map. Ap-

pendix Appendix C lists the resulting threats and links each entry to one or more argument IDs; the full Toulmin decompositions of these arguments are provided in Appendix B. For each threat, we record the affected layer (e.g., WebPKI/TLS, qualified identity plane, trust-list governance, user-agent UI), the main preconditions/assumptions, the plausible impact, and mitigation or design levels (descriptive, not prescriptive). This threat model represents a synthesis of risks and failure modes discussed by stakeholders and reflected in the standards and regulatory setting. It is not presented as the discovery of new attacks, nor as an estimate of likelihood. Its purpose is to make explicit which security outcomes depend on governance and implementation choices (e.g., 1-QWAC vs. 2-QWAC consumption, trust-list update and distrust procedures, revocation/status handling, and how qualified identity is displayed).

6.4. Limitations and Future Work

This study analyzes the QWAC controversy as a socio-technical dispute in which many claims are conditional on implementation and enforcement choices that are still evolving. Consequently, our results should be interpreted as a structured representation of stakeholder argumentation (claims, grounds, warrants, qualifiers, rebuttals) rather than as outcome-level evidence about security effectiveness. As such, we face several limitations: (1) The evidence base is documentary and qualitative: we do not observe large-scale deployment outcomes, controlled experiments, or longitudinal behavioral effects, which constrains external validity for user- and incident-rate claims. (2) The controversy is temporally sensitive: some prominent arguments were shaped by draft text, and both delegated acts and technical specifications continue to evolve. We aimed to mitigate these limitation – to the best of our ability – by (1) grounding claims in a vast multivocal corpus that includes normative texts (law/standards) and operational governance artifacts (root program and interoperability documents), (2) interpreting arguments against the adopted Regulation (EU) 2024/1183 and ETSI TS 119 411-5 (v2.1.1) unless stated otherwise, and (3) using Toulmin-structured coding to make assumptions and conditions explicit so that readers can assess where arguments depend on unsettled premises. To further operationalize the security-related parts of the controversy without implying outcome validation, we include an appendix threat model that consolidates recurring cross-layer failure modes and associated design levers.

Correspondingly, future empirical work is still needed to evaluate how qualified-identity presentation impacts user behavior and fraud outcomes in realistic settings, to measure issuance, adoption, and incident dynamics as implementations mature (including revocation/distrust latency across governance layers). Moreover, the interviews suggested a strong need to compare supervisory practice across Member States and model heterogeneity in compliance and enforcement, as well as develop reproducible conformance suites and reference implementations for interoperability-critical components (trust-list processing, qualified identity parsing, and 2-QWAC binding and status behaviors).

7. Conclusion

Our evaluation revealed three main takeaways. First, arguments concerning security and trust focus on whether mandatory recognition and extended identity disclosure provide measurable security benefits for end users or introduce additional attack vectors. Second, governance and authority arguments revolve around the balance between public accountability and reduced influence from large technology companies against the risks of scope creep, politicization, and unclear legal remedies within the EU. Finally, concerns about compliance and interoperability are conditional, with coexistence between QWACs and the WebPKI appearing feasible only if display, validation, and revocation behaviors are consistently defined and implemented.

The controversy surrounding QWACs highlights the challenge of aligning EU regulatory ambitions with global WebPKI practices. Most arguments proved to be dependent on specific conditions, underscoring that the effectiveness of QWACs relies not only on their formal definition but also largely on their practical implementation. This includes whether browsers adopt truly user-friendly and transparent indicators, whether supervisory bodies ensure timely escalation and oversight, and whether mechanisms such as CT are fully integrated. In this sense, QWACs serve as a proxy for broader struggles over who governs trust on the web.

From our point of view, the QWAC controversy also suggests that the increasing entanglement of legal and technical dimensions of security warrants closer attention within academic discourse. Our structured examination of the arguments articulated in the context of the QWAC debate indicates that not only industry contributions – whether from supporters or critics of eIDAS Article 45 – but also the associated open letters tended to

reflect a limited degree of nuance. Particularly our interviews with signatories of one or both open letters revealed that some were not familiar with important regulatory context that may have led to more differentiated assessments, and participants generally acknowledged the relevance of several considerations advanced by proponents. Moreover, discussions held during interviews and at industry conferences suggested that the tone of the open letters was perceived as less balanced than stakeholders would typically expect from scholarly engagement. Against this background, we propose that future debates situated at the interface of legal and technical security mechanisms would benefit from greater epistemic caution. Even when motivated by well-intentioned objectives, researchers with limited exposure to the legal context might consider adopting weaker formulations to mitigate the risk that categorical statements are instrumentalised by interested parties, thereby helping to preserve the credibility of academic contributions in both public and private arenas. Without sustaining such legitimacy, there is a pronounced risk that diminished credibility spills over into policy discussions that appear considerably less contested within the scientific community – such as those concerning bans of or backdoors on end-to-end encryption – thereby making policy-makers less inclined to rely on academic expertise precisely in situations where critical guidance is most needed.

In summary, our study contributes a consolidated stakeholder argument map and catalogue (Appendix A.2) that disentangles the QWAC controversy by integrating systematic evidence, stakeholder interviews, and the temporal development of the debate, and translates these results into an argument-linked threat model (Table C.1) for implementation-oriented reasoning. The results clarify where the controversy is driven by governance assumptions, where it is driven by usability/consumption assumptions, and where it is driven by interoperability and enforcement assumptions. Because many claims remain conditional on implementation and supervision, this paper should be read as mapping the dispute and its dependencies rather than resolving it through outcome-level measurements. Accordingly, QWACs are neither a comprehensive solution to website authentication nor an intrinsic source of risk; they are regulatory instruments whose practical effects depend on the interaction between legal mandates, technical standards, and operational enforcement.

Ethics considerations

Ethical approval was obtained from the University of Luxembourg's Ethics Review Panel, with oversight from the Data Protection Office to ensure compliance with the GDPR and national law. Because the study employed interviews that collected personal data, we implemented an information sheet and obtained informed consent to explain the purpose, data use, storage, and participants' rights (including access, correction, and withdrawal). Participation was voluntary. Data are securely stored with restricted access and defined retention. Participants could withdraw at any time without penalty and were pseudonymized in this study. These measures minimized risks, protected confidentiality, and maintained legal and institutional compliance.

Acknowledgments

This research was supported in part by Luxembourg's Ministry for Digitalisation, PayPal, and the Luxembourg National Research Fund (FNR) (PEARL grant reference FNR13342933, PABLO grant reference FNR163267543, and NCER-FT grant reference FNR165704683), as well as by the Ministry of Finance of Luxembourg through the FutureFinTech National Centre of Excellence in Research & Innovation. For open-access purposes, the authors have applied a CC BY 4.0 license to any Author Accepted Manuscript arising from this submission.

Declaration of AI Usage

During the preparation of this manuscript, the authors used Grammarly's generative AI assistant and OpenAI's ChatGPT 5 (Pro) to enhance language clarity and readability and convert research notes into structured text. After using these tools, the authors thoroughly reviewed, edited, and validated all content and took full responsibility for the final version of the publication. These tools have not been used to process any personal or copyright data; instead, they have been used to process materials owned by the authors and publicly available documents.

References

- Aas, J., Barnes, R., Case, B., Durumeric, Z., Eckersley, P., Flores-López, A., Halderman, J.A., Hoffman-Andrews, J., Kasten, J., Rescorla, E., others, 2019. Let's Encrypt: An Automated Certificate Authority to Encrypt the Entire Web, in: Proceedings of the ACM SIGSAC Conference on Computer and Communications Security, pp. 2473–2487. doi:10.1145/3319535.3363192.
- Adriano Santoni, 2025. Are all QWACs the same? URL: <https://www.linkedin.com/pulse/all-qwacs-same-adriano-santoni-9snqf/>.
- Amann, J., Gasser, O., Scheitle, Q., Brent, L., Carle, G., Holz, R., 2017. Mission Accomplished?: HTTPS Security After DigiNotar, in: Proceedings of the 2017 Internet Measurement Conference, ACM, London United Kingdom. pp. 325–340. doi:10.1145/3131365.3131401.
- Bailey, C., 2022. eIDAS 2, ARTICLE 45: Where we are, and possible outcomes. URL: <https://www.enisa.europa.eu/events/trust-services-forum-ca-day-2022/presentations/chris-bailey-enisa-trust-services-forum-2022.pdf>.
- Barbureau, T., Weigl, L., Pocher, N., 2024. Financial Regulation, Political Context, and Technology in the European Union, in: Fridgen, G., Guggenberger, T., Sedlmeir, J., Urbach, N. (Eds.), Decentralization Technologies. Springer, pp. 19–46. URL: https://link.springer.com/10.1007/978-3-031-66047-4_2, doi:10.1007/978-3-031-66047-4_2. series Title: Financial Innovation and Technology.
- Berbecaru, D.G., Lioy, A., 2023. An Evaluation of X.509 Certificate Revocation and Related Privacy Issues in the Web PKI Ecosystem. IEEE Access 11, 79156–79175. doi:10.1109/ACCESS.2023.3299357.
- Biddle, R., Van Oorschot, P.C., Patrick, A.S., Sobey, J., Whalen, T., 2009. Browser Interfaces and Extended Validation SSL Certificates: An Empirical Study, ACM. pp. 19–30. doi:10.1145/1655008.1655012.
- Bitkom, 2020. Position Paper: Independence of Trust Services Providers from Browser and Operating Systems. URL: https://bitkom.org/sites/default/files/2020-09/20200918_bitkom-position-independence-of-trust-service-providers-final.pdf.

- Bitkom, 2022. Websitezertifikate zur Stärkung der europäischen Souveränität. URL: https://www.bitkom.org/sites/main/files/2022-05/20220320_Bitkom_Position_QWACs.pdf.
- Boeyen, S., Santesson, S., Polk, T., Housley, R., Farrell, S., Cooper, D., 2008. RFC 5280: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. Request for Comments RFC 5280. Internet Engineering Task Force. URL: <https://datatracker.ietf.org/doc/rfc5280>.
- Braun, J., Volk, F., Classen, J., Buchmann, J., Mühlhäuser, M., 2014. CA Trust Management for the Web PKI. *Journal of Computer Security* 22, 913–959. doi:10.3233/JCS-140509.
- Bundesdruckerei, 2022. Qualified Website Authentication Certificates: EU Plan to Boost Web Security. URL: <https://www.bundesdruckerei.de/en/innovation-hub/eu-plan-boost-web-security>.
- Bundeskartellamt, 2022. Proceedings in Transport Layer Security Certificates against Google terminated. URL: https://www.bundeskartellamt.de/SharedDocs/Meldung/EN/AktuelleMeldungen/2022/21_01_2022_Fallbericht_Google_Zertifikate.html.
- CA/B Forum, 2020. Proposed Change to the Technical Profile for Qualified Website Authentication Certificates (QWACs). URL: <https://archive.cabforum.org/pipermail/servercert-wg/attachments/20200114/3a5fa74c/attachment-0001.pdf>.
- CA/B Forum, 2025. Baseline Requirements for the Issuance and Management of Publicly-Trusted TLS Server Certificates. Server Cert WG Version 2.1.7. URL: <https://cabforum.org/working-groups/server/baseline-requirements/documents/>.
- Campbell, S., Greenwood, M., Prior, S., Shearer, T., Walkem, K., Young, S., Bywaters, D., Walker, K., 2020. Purposive sampling: complex or simple? Research case examples. *Journal of Research in Nursing* 25, 652–661. URL: <https://journals.sagepub.com/doi/10.1177/1744987120927206>, doi:10.1177/1744987120927206.

- CCADB, 2018. Trust Service Provider Technical Best Practices Considering the EU eIDAS Regulation (910/2014). URL: https://www.ccadb.org/documents/TSP_Technical_Best_Practices_eIDAS.pdf.
- CCADB, 2020a. Executive Summary: Interoperability Challenges of QWACs in Binding to Connections instead of Domains. Technical Report. URL: https://www.ccadb.org/documents/qualified_website_authentication_certificates_interoperability.pdf.
- CCADB, 2020b. Qualified Website Authentication Certificates (QWACs) Interoperability. URL: https://www.ccadb.org/documents/Qualified_Website_Authentication_Certificates_Interoperability.pdf.
- Certificate Transparency, 2020. How CT Fits Into the Wider Web PKI Ecosystem. URL: <https://certificate.transparency.dev/>.
- certSIGN, 2024. Repository of Certification Policies, Certification Practice Statements, and Other Relevant Documents for Our Trusted Services. URL: <https://www.certsign.ro/en/repository/>.
- Chuat, L., Abdou, A., Sasse, R., Sprenger, C., Basin, D., Perrig, A., 2020. SoK: Delegation and Revocation, the Missing Links in the Web's Chain of Trust, in: 2020 IEEE European Symposium on Security and Privacy (EuroS&P), pp. 624–638. URL: <https://ieeexplore.ieee.org/abstract/document/9230363>, doi:10.1109/EuroSP48549.2020.00046.
- Claburn, T., 2023. Bad eIDAS: Europe Ready to Intercept, Spy on Your Encrypted HTTPS Connections. URL: https://www.theregister.com/2023/11/08/europe_eidas_browser/.
- Cybersecurity Experts, 2022. Open Letter: Joint statement of cybersecurity experts on the EU's proposed eIDAS reform. URL: https://epicenter.works/fileadmin/user_upload/eIDAS_Open_Letter-2023-11-01-Academics_NGOs.pdf.
- Degen, K., Teubner, T., 2024. Wallet wars or digital public infrastructure? Orchestrating a digital identity data ecosystem from a government perspective. *Electronic Markets* 34, 50. URL: <https://doi.org/10.1007/s12525-024-00731-1>, doi:10.1007/s12525-024-00731-1.

- Delignat-Lavaud, A., Abadi, M., Birrell, A., Mironov, I., Wobber, T., Xie, Y., 2014. Web PKI: Closing the Gap Between Guidelines and Practices, in: Network and Distributed System Security Symposium, USENIX Association. URL: <https://users.soe.ucsc.edu/~abadi/Papers/ndss14.pdf>.
- Drury, V., Meyer, U., 2019. Certified Phishing: Taking a Look at Public Key Certificates of Phishing Websites, in: Proceedings of the 15th USENIX Conference on Usable Privacy and Security, USENIX Association.
- Eastlake 3rd, D.E., 2011. RFC 6066: Transport Layer Security (TLS) Extensions: Extension Definitions. Request for Comments RFC 6066. Internet Engineering Task Force. URL: <https://datatracker.ietf.org/doc/rfc6066>, doi:10.17487/RFC6066.
- Electronic Frontier Foundation, 2023. Article 45 Will Roll Back Web Security by 12 Years. URL: <https://www.eff.org/deeplinks/2023/11/article-45-will-roll-back-web-security-12-years>.
- Electronic Frontier Foundation (EFF), 2022. What the Duck? Why an EU Proposal to Require "QWACs" Will Hurt Internet Security. URL: <https://www.eff.org/deeplinks/2022/02/what-duck-why-eu-proposal-require-qwacs-will-hurt-internet-security>.
- ENISA, 2016. Qualified Website Authentication Certificates: Promoting Consumer Trust in the Website Authentication Market. Technical Report. URL: <https://doi.org/10.2824/464966>.
- ENISA, 2017. Security guidelines on the appropriate use of qualified website authentication certificates: Guidance for users. Technical Report Version 2.0. URL: <https://www.enisa.europa.eu/publications/security-guidelines-on-the-appropriate-use-of-qualified-website-authentication-certificates>.
- Entrust Corporation, 2021. Designing the New eIDAS 2 Browser UI. URL: https://cabforum.org/uploads/05_chris-bailey_20210900-ca-day-designing-the-new-eidas-2-browser-ui_1_.pdf.
- Entschew, E., van Brouwershaven, P., 2024. QWACs in the Context of the Trust Spaces of the Browsers and eIDAS 2.0. *Datenschutz und Datensicherheit* 48, 246–250. doi:10.1007/s11623-024-1918-x.

- Entschew, E., Hall, K., Bailey, C., Nguyen, K., 2022. A New eIDAS Beginning for QWACs. *Datenschutz und Datensicherheit* 46, 217–224. doi:10.1007/s11623-022-1591-x.
- ESD, 2022a. Mozilla Website Pushes Serious eIDAS Misinformation to Political Decision Makers and Public. Technical Report. URL: https://www.linkedin.com/posts/european-signature-dialog_mozilla-campaign-pushes-serious-misinformation-activity-6978078620279824384-ByAc/.
- ESD, 2022b. Position Paper on the Provisions of Article 45 (QWACs). Technical Report. URL: [https://www.european-signature-dialog.eu/Position_Paper_on_the_provisions_of_Article_45_\(QWACs\).pdf](https://www.european-signature-dialog.eu/Position_Paper_on_the_provisions_of_Article_45_(QWACs).pdf).
- ESD, 2022c. Vote for a strong sovereign Europe that can guarantee data protection rights to its citizens and control the EU Trust Space – eIDAS Art. 45. Technical Report. URL: https://www.european-signature-dialog.eu/Art_45_eIDAS_arguments_in_favor.pdf.
- ESD, 2023a. eIDAS Art. 45.2 Requires Browsers to Show EU Citizens a “User Friendly UI”. Why Is This Important? Technical Report. URL: <https://www.european-signature-dialog.eu/ourimpact/>.
- ESD, 2023b. ESD Experts Support Trilogue Compromise and Emphasize Necessity for Highest Security of the Internet. Technical Report. URL: https://www.european-signature-dialog.eu/ESD_experts_support_trilogue_Art.45_results-6nov2023.pdf.
- ESD, 2023c. Five Top Reasons for Adoption of the New eIDAS Article 45 on QWACs. Technical Report. URL: https://www.european-signature-dialog.eu/5reasons_for_adoption_Art45_QWACs.pdf.
- ETSI, 2023a. EN 319 411-1 (1.4.1): Electronic Signatures and Infrastructures (ESI); Policy and security requirements for Trust Service Providers issuing certificates; Part 1: General requirements. Technical Report 1.4.1. URL: https://portal.etsi.org/webapp/workprogram/Report_WorkItem.asp?WKI_ID=64273.
- ETSI, 2023b. EN 319 411-2 (2.5.1): Electronic Signatures and Infrastructures (ESI); Policy and security requirements for Trust Service Providers issuing

- certificates; Part 2: Requirements for trust service providers issuing EU qualified certificates. Technical Report 2.5.1. URL: https://portal.etsi.org/webapp/workprogram/Report_WorkItem.asp?WKI_ID=67516.
- ETSI, 2023c. EN 319 412-1 (1.5.1): Electronic Signatures and Infrastructures (ESI); Certificate Profiles; Part 1: Overview and Common Data Structures. 1.5.1. URL: https://portal.etsi.org/webapp/WorkProgram/Report_WorkItem.asp?WKI_ID=64271.
- ETSI, 2023d. EN 319 412-3 (1.3.1): Electronic Signatures and Infrastructures (ESI); Certificate Profiles; Part 3: Certificate Profile for Certificates Issued to Legal Persons. Technical Report. URL: https://portal.etsi.org/webapp/WorkProgram/Report_WorkItem.asp?WKI_ID=68498.
- ETSI, 2023e. EN 319 412-5 (2.4.1): Electronic Signatures and Infrastructures (ESI); Certificate Profiles; Part 5: QCStatements. Technical Report. URL: https://portal.etsi.org/webapp/WorkProgram/Report_WorkItem.asp?WKI_ID=68497.
- ETSI, 2023f. TR 119 476 (v1.1.1): Electronic Signatures and Infrastructures (ESI); Analysis of selective disclosure and zero-knowledge proofs applied to Electronic Attestation of Attributes. Technical Report 1.1.1. URL: https://portal.etsi.org/webapp/workprogram/Report_WorkItem.asp?WKI_ID=67975.
- ETSI, 2023g. TS 119 411-6 (1.1.1): Electronic Signatures and Infrastructures (ESI); Policy and security requirements for Trust Service Providers issuing certificates; Part 6: Requirements for Trust Service Providers issuing publicly trusted S/MIME certificates. Technical Report 1.1.1. URL: https://portal.etsi.org/webapp/workprogram/Report_WorkItem.asp?WKI_ID=67990.
- ETSI, 2024a. EN 319 401 (v3.1.1): Electronic Signatures and Trust Infrastructures (ESI); General Policy Requirements for Trust Service Providers. Technical Report 3.1.1. URL: https://portal.etsi.org/webapp/workprogram/Report_WorkItem.asp?WKI_ID=67940.
- ETSI, 2024b. TR 119 411-4 (1.2.1): Electronic Signatures and Infrastructures (ESI); Policy and security requirements for Trust Service Providers issuing certificates; Part 4: Checklist supporting audit of TSP

against ETSI EN 319 411-1 or ETSI EN 319 411-2. Technical Report 1.2.1. URL: https://portal.etsi.org/webapp/workprogram/Report_WorkItem.asp?WKI_ID=68503.

ETSI, 2025. TS 119 411-5 V2.1.1: Electronic Signatures and Trust Infrastructures (ESI): Policy and security requirements for Trust Service Providers issuing certificates – Part 5: Implementation of qualified certificates for website authentication as in amended Regulation 910/2014. Technical Report TS 119 411-5. ETSI. URL: https://www.etsi.org/deliver/etsi_ts/119400_119499/11941105/02.01.01_60/ts_11941105v020101p.pdf, doi:RTS/ESI00194115v211. version Number: 2.1.1.

European Commission (Ed.), 2015. Joint practical guide of the European Parliament, the Council and the Commission for persons involved in the drafting of European Union legislation. Second edition ed., Publications Office, Luxembourg. URL: <https://eur-lex.europa.eu/content/techleg/KB0213228ENN.pdf>, doi:10.2880/89965.

European Commission, 2020a. CEF eSignature Pilot on ntQWACs. URL: https://ec.europa.eu/futurium/sites/futurium/files/ntqwac_pilot.pdf.

European Commission, 2020b. Evaluation Study of the Regulation no. 910/2014 (eIDAS Regulation). Technical Report SMART 2019/0046. URL: <https://digital-strategy.ec.europa.eu/en/library/evaluation-study-regulation-no9102014-eidas-regulation>.

European Commission, 2021a. eIDAS Dashboard: What is a QWAC? URL: <https://eidas.ec.europa.eu/efda/discover/qwac>.

European Commission, 2021b. Impact Assessment Report. Technical Report. URL: <https://digital-strategy.ec.europa.eu/en/library/trusted-and-secure-european-e-id-regulation>.

European Commission, 2021c. Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL amending Regulation (EU) No 910/2014 as regards establishing a framework for a European Digital Identity. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM%3A2021%3A281%3AFIN>.

- European Commission, 2021d. Report From the Commission to the European Parliament and the Council on the Evaluation of Regulation (EU) No 910/2014 on Electronic Identification and Trust Services for Electronic Transactions in the Internal Market (eIDAS). Technical Report. URL: <https://digital-strategy.ec.europa.eu/en/library/trusted-and-secure-european-e-id-evaluation-regulation>.
- European Commission, 2021e. Study to Support the Impact Assessment for the Revision of the eIDAS Regulation. Technical Report 2020/666. European Commission. URL: <https://digital-strategy.ec.europa.eu/en/library/study-support-impact-assessment-revision-eidas-regulation>.
- European Commission, 2022. European Digital Identity. URL: https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/european-digital-identity_en.
- European Commission, 2023a. eIDAS Dashboard: Trust service providers: QWAC. URL: <https://eidas.ec.europa.eu/efda/trust-services/browse/eidas/tls/search/type>.
- European Commission, 2023b. European Digital Identity – Questions and Answers. URL: https://ec.europa.eu/commission/presscorner/detail/en/QANDA_21_2664.
- European Commission, 2023c. Final Agreement on EU Digital Identity Wallet. URL: https://ec.europa.eu/commission/presscorner/detail/en/IP_23_5651.
- European Commission, 2024. European Digital Identity Wallet Architecture and Reference Framework. Technical Report. URL: <https://github.com/eu-digital-identity-wallet/eudi-doc-architecture-and-reference-framework/blob/main/docs/arf.md>.
- European Parliament, 2024. The ubiquitous digital single market - Fact Sheets on the European Union. URL: <https://www.europarl.europa.eu/factsheets/en/sheet/43/the-ubiquitous-digital-single-market>.
- European Parliament, European Council, 2014. Regulation (EU) No 910/2014 on electronic identification and trust services for electronic

- transactions in the internal market and repealing Directive 1999/93/EC. URL: <http://data.europa.eu/eli/reg/2014/910/oj/eng>.
- European Union, 2014. Regulation (eu) no 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. URL: <http://data.europa.eu/eli/reg/2014/910/oj/eng>.
- European Union, 2015. Directive (eu) 2015/2366 of the european parliament and of the council of 25 november 2015 on payment services in the internal market, amending directives 2002/65/ec, 2009/110/ec and 2013/36/eu and regulation (eu) no 1093/2010, and repealing directive 2007/64/ec. URL: <http://data.europa.eu/eli/dir/2015/2366/oj/eng>.
- European Union, 2016. Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/ec (general data protection regulation). URL: <http://data.europa.eu/eli/reg/2016/679/oj/eng>.
- European Union, 2024. Regulation (eu) 2024/1183 of the european parliament and of the council of 11 april 2024 amending regulation (eu) no 910/2014 as regards establishing the european digital identity framework. URL: <http://data.europa.eu/eli/reg/2024/1183/oj/eng>.
- Feisty Duck, 2019. Firefox and Chrome Will Remove GUI Indicator for Extended Validation Certificates. URL: https://www.feistyduck.com/newsletter/issue_56_firefox_and_chrome_will_remove_gui_indicator_for_extended_validation_certificates.
- Feisty Duck, 2022. EU Plans to Mandate Less Secure Certificates in Browsers. URL: https://www.feistyduck.com/newsletter/issue_86_eu_plans_to_mandate_less_secure_certificates_in_browsers.
- Felt, A.P., Reeder, R.W., Ainslie, A., Harris, H., Walker, M., Thompson, C., Acer, M.E., Morant, E., Consolvo, S., 2016. Rethinking Connection Security Indicators, in: Proceedings of the 12th Symposium on Usable Privacy and Security, USENIX Association.

- Fiedler, A., Thiel, C., 2014. The need of European White Knights for the TLS/SSL Certificate System, in: Highlights of the Information Security Solutions Europe Conference: Securing Electronic Business Processes, Springer. pp. 170–174. doi:10.1007/978-3-658-06708-3_13.
- Garousi, V., Felderer, M., Mäntylä, M.V., 2016. The Need for Multivocal Literature Reviews in Software Engineering: Complementing Systematic Literature Reviews With Grey Literature, in: Proceedings of the 20th International Conference on Evaluation and Assessment in Software Engineering, ACM. URL: <https://dl.acm.org/doi/10.1145/2915970.2916008>, doi:10.1145/2915970.2916008.
- Google, 2025. HTTPS Encryption on the Web. URL: <https://transparencyreport.google.com/https/overview>.
- Gregor, S., 2006. The Nature of Theory in Information Systems1. MIS Quarterly 30, 611–642. URL: <https://misq.umn.edu/misq/article/30/3/611/419/The-Nature-of-Theory-in-Information-Systems1>, doi:10.2307/25148742.
- Grindal, K., Mueller, M., Srivastava, V., 2025. Non-Governmental Governance of Trust on the Internet: WebPKI as Public Good. Journal of Cybersecurity 11, tyaf018. URL: <https://academic.oup.com/cybersecurity/article/doi/10.1093/cybsec/tyaf018/8225340>, doi:10.1093/cybsec/tyaf018.
- Heidebrecht, S., 2024. From Market Liberalism to Public Intervention: Digital Sovereignty and Changing European Union Digital Single Market Governance. JCMS: Journal of Common Market Studies 62, 205–223. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/jcms.13488>, doi:10.1111/jcms.13488. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/jcms.13488>.
- Helme, S., 2021. Top 1 Million Analysis - November 2021. Technical Report. URL: <https://scotthelme.co.uk/top-1-million-analysis-november-2021>.
- Helme, S., 2023. What the QWAC?! URL: <https://scotthelme.co.uk/what-the-qwac/>.

- Hunt, T., 2019. Extended Validation Certificates are (Really, Really) Dead. URL: <https://www.troyhunt.com/extended-validation-certificates-are-really-really-dead/>.
- Identityum, 2022. EU vs USA OR EU with USA: REVISING THE EIDAS 2.0 REGULATION. URL: <https://identityum.com/eu-vs-usa-or-eu-with-usa-revising-the-eidas-2-0-regulation/>.
- Information Security Media Group, 2023. Browser Makers and EU Face Off Over QWACs. URL: <https://www.inforisktoday.com/browser-makers-eu-face-off-over-qwacs-a-21850>.
- Internet Society, Center for Democracy & Technology, Electronic Frontier Foundation, Epicenter.works, 2023. Civil Society Experts Voice Concern as New EU Digital Identity Regulation Finalized. URL: <https://www.internetsociety.org/resources/doc/2023/qualified-web-authentication-certificates-qwacs-in-eidas/>.
- Jackson, C., Simon, D.R., Tan, D.S., Barth, A., 2007. An Evaluation of Extended Validation and Picture-In-Picture Phishing Attacks, in: International Conference on Financial Cryptography and Data Security, pp. 281–293.
- Kallio, H., Pietilä, A.M., Johnson, M., Kangasniemi, M., 2016. Systematic Methodological Review: Developing a Framework for a Qualitative Semi-Structured Interview Guide. *Journal of Advanced Nursing* 72, 2954–2965. doi:10.1111/jan.13031.
- Keizer, G., 2011. Hackers Spied on 300,000 Iranians Using Fake Google Certificate. URL: <https://www.computerworld.com/article/1545206/hackers-spied-on-300-000-iranians-using-fake-google-certificate.html>.
- Klimas, T., Vaiciukaite, J., 2008. The Law of Recitals in European Community Legislation. *ILSA Journal of International and Comparative Law* URL: <https://www.semanticscholar.org/paper/The-Law-of-Recitals-in-European-Community-Klimas-Vaiciukaite/316a3e2169e46a86d37383a26eebca7b23ec0eba>.
- Koschuch, M., Wagner, R., 2015. Trust Revoked — Practical Evaluation of OCSP- and CRL-Checking Implementations, in: Obaidat, M.S., Holzinger,

- A., Filipe, J. (Eds.), E-Business and Telecommunications. volume 554, pp. 26–33. URL: http://link.springer.com/10.1007/978-3-319-25915-4_2, doi:10.1007/978-3-319-25915-4_2. series Title: Communications in Computer and Information Science.
- Kudra, A., Seegebart, C., Schwalm, S., 2022. Ein digitaler Vertrauensraum für Identitäten und Dienste – Europa ist auf dem richtigen Weg: Ein Impuls. Datenschutz und Datensicherheit - DuD 46, 9–11. URL: <https://link.springer.com/10.1007/s11623-022-1552-4>, doi:10.1007/s11623-022-1552-4.
- Kuzman, T., Mozetič, I., Ljubešić, N., 2023. Automatic Genre Identification for Robust Enrichment of Massive Text Collections: Investigation of Classification Methods in the Era of Large Language Models. Machine Learning and Knowledge Extraction 5, 1149–1175. URL: <https://www.mdpi.com/2504-4990/5/3/59>, doi:10.3390/make5030059.
- Laurie, B., 2014. Certificate Transparency. Communications of the ACM 57, 40–46. doi:10.1145/2659897.
- Laurie, B., Langley, A., Kasper, E., Messeri, E., Stradling, R., 2021. RFC 9162: Certificate Transparency Version 2.0. Request for Comments RFC 9162. Internet Engineering Task Force. URL: <https://datatracker.ietf.org/doc/rfc9162>.
- Levy, Y., Ellis, T.J., 2006. A Systems Approach to Conduct an Effective Literature Review in Support of Information Systems Research. Informing Science: The International Journal of an Emerging Transdiscipline 9, 181–212. doi:10.28945/479.
- Liu, D., Xiong, M., 2024. Keeping balance between loyalty and modification: a Toulminian model as analytical framework. Humanities and Social Sciences Communications 11, 639. URL: <https://www.nature.com/articles/s41599-024-03151-w>, doi:10.1057/s41599-024-03151-w.
- LuxTrust S.A., 2024. QSealC & QWAC Certificates for PSD2. URL: <https://www.luxtrust.com/en/professionals/our-digital-solutions/meet-psd2-requirements>.

- Luzi, D., 2000. Trends and Evolution in the Development of Grey Literature: A Review. *International Journal on Grey Literature* 1, 106–117. doi:10.1108/14666180010345537.
- Martius, K., Hühnlein, T., Hühnlein, D., Wich, T., 2024. Trustworthy QWACs – Fact or Fiction?, in: *Proceedings of the Open Identity Summit*, Gesellschaft für Informatik e.V.. pp. 189–194. doi:10.18420/OID2024_18.
- van der Meulen, N., . DigiNotar: Dissecting the First Dutch Digital Disaster. *Journal of Strategic Security* 6, 46–58. URL: <https://www.jstor.org/stable/26466760>.
- Moher, D., Liberati, A., Tetzlaff, J., Altman, D.G., 2009. Preferred Reporting Items for Systematic Reviews and Meta-Analyses: The PRISMA Statement. *British Medical Journal* 339. doi:10.1136/bmj.b2535.
- Mozilla, 2020. The EU’s Current Approach to QWACs will Undermine Security on the Open Web. URL: <https://blog.mozilla.org/netpolicy/2020/10/08/the-eus-current-approach-to-qwacs-qualified-website-authentication-certificates-will-undermine-security-on-the-open-web>.
- Mozilla, 2021. Position paper on the European Commission’s legislative proposal to revise the eIDAS Regulation.
- Mozilla, 2022. Discover How QWACs Can Put You at Risk. URL: <https://securityriskahead.eu/>.
- Mozilla, 2023a. How does Article 45 enable the interception of web traffic? URL: <https://last-chance-for-eidas.org/art45interception.html>.
- Mozilla, 2023b. Last Chance to fix eIDAS: Secret EU law Threatens Internet Security. URL: <https://last-chance-for-eidas.org/>.
- Mozilla, 2023c. Root Store Policy. URL: <https://www.mozilla.org/en-US/about/governance/policies/security-group/certs/policy/>. version 2.9.
- Mozilla, 2024. EUROPEAN PARLIAMENT VOTES ON eIDAS, AVERTING THREAT TO WEB SECURITY. URL: <https://>

[//securityriskahead.eu/wp-content/uploads/2024/02/Mozilla_press-release_eIDAS-EP-Plenary-adoption.pdf](https://securityriskahead.eu/wp-content/uploads/2024/02/Mozilla_press-release_eIDAS-EP-Plenary-adoption.pdf).

- Mozilla, Akamai, Bytecode Alliance, Cisco, Cloudflare, DNS0.EU, Fastly, Internet Security Research Group, Linux Foundation, Mullvad, OpenSSF, Sigstore, 2023. Industry Joint Statement on Article 45 in the EU’s eIDAS Regulation. URL: <https://blog.mozilla.org/netpolicy/files/2023/11/eIDAS-Industry-Letter.pdf>.
- Muffett, A., 2023a. Hot on the Heels of ChatControl and in the Name of “Identity” and “Consumer Choice” the EU Seeks the Ability to Undetectably Spy on HTTPS Communication. URL: <https://alecmuffett.com/article/108139>.
- Muffett, A., 2023b. Leaked Document From MEPs Lays Out a Delusional, Paranoid Joint (Mis)understanding of how Web Standards Work, Makes Demands for Government, not Experts, to Dictate how Web Works / for HTTPS Backdoor. URL: <https://alecmuffett.com/article/108519>.
- Orlikowski, W.J., 1992. The Duality of Technology: Rethinking the Concept of Technology in Organizations. *Organization Science* 3, 398–427. URL: <https://pubsonline.informs.org/doi/10.1287/orsc.3.3.398>, doi:10.1287/orsc.3.3.398.
- Raman, R.S., Evdokimov, L., Wurstrow, E., Halderman, J.A., Ensafi, R., 2020. Investigating Large Scale HTTPS Interception in Kazakhstan, in: *Proceedings of the ACM Internet Measurement Conference*, pp. 125–132. doi:10.1145/3419394.3423665.
- Recorded Future News, 2023. EU Urged to Drop New Law That Could Allow Member States to Intercept and Decrypt Global Web Traffic. *The Record* URL: <https://therecord.media/eu-urged-to-drop-law-website-authentication-certificates>.
- Rescorla, E., 2018. RFC 8446: The Transport Layer Security (TLS) Protocol Version 1.3. Request for Comments RFC 8446. Internet Engineering Task Force. URL: <https://datatracker.ietf.org/doc/rfc8446>.
- Rescorla, E., 2022. Can we agree on the facts about QWACs? URL: <https://educatedguesswork.org/posts/eidas-article45/>.

- Rescorla, E., Dierks, T., 2008. RFC 5246: The Transport Layer Security (TLS) Protocol Version 1.2. Request for Comments RFC 5246. Internet Engineering Task Force. URL: <https://datatracker.ietf.org/doc/rfc5246>.
- Saldaña, J., 2021. The Coding Manual for Qualitative Researchers. 4e ed., SAGE, Los Angeles London New Delhi Singapore Washington DC Melbourne.
- Scheitle, Q., Gasser, O., Nolte, T., Amann, J., Brent, L., Carle, G., Holz, R., Schmidt, T.C., Wählisch, M., 2018. The Rise of Certificate Transparency and its Implications on the Internet Ecosystem, in: Proceedings of the Internet Measurement Conference, pp. 343–349. doi:10.1145/3278532.3278562.
- Schultze, U., Avital, M., 2011. Designing interviews to generate rich data for information systems research. *Information and Organization* 21, 1–16. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1471772710000412>, doi:10.1016/j.infoandorg.2010.11.001.
- Schwalm, S., 2023. QWAC or not QWAC is That the Question? URL: <https://medium.com/@schwalm.steffen/qwac-or-not-qwac-is-that-the-question-80b7a145db9d>.
- Scientists and NGOs, 2023a. Open Letter: Joint statement of scientists and NGOs on the EU’s proposed eIDAS reform. URL: https://epicenter.works/fileadmin/user_upload/eIDAS_Open_Letter-2023-11-01-Academics_NGOs.pdf.
- Scientists and NGOs, 2023b. Response Briefing Note on the Discussion Around Qualified Website Authentication Certificates (QWACs) - Art. 45 of eIDAS Regulation. URL: https://homes.esat.kuleuven.be/~preneel/response_to_eIDAS_briefing_note_1december2023.pdf.
- Sectigo Limited, 2024. Qualified Website Authentication Certificates. URL: <https://www.sectigo.com/ssl-certificates-tls/qwac-businesses>.
- SSLInsights, 2025. Latest SSL Certificates Statistics 2025. URL: <https://sslinsights.com/ssl-certificates-statistics/>.

- Stojmenovic, M., Biddle, R., 2018. Who are They? Website Authentication: Certificates and Identity, in: Who Are You Workshop at USENIX Symposium on Usable Privacy and Security. URL: <https://wayworkshop.org/2018/papers/way2018-stojmenovic.pdf>.
- Tehrani, P.F., Osterweil, E., Schmidt, T.C., Wählisch, M., 2024. How to Measure TLS, X.509 Certificates, and Web PKI: A Tutorial and Brief Survey. URL: <https://arxiv.org/abs/2401.18053>. arXiv:2401.18053 [cs].
- Thompson, C., Shelton, M., Stark, E., Walker, M., Schechter, E., Felt, A., 2019. The Web's Identity Crisis: Understanding the Effectiveness of Website Identity Indicators, in: Proceedings of the 28th USENIX Conference on Security Symposium, pp. 1715–1732. URL: <https://www.usenix.org/conference/usenixsecurity19/presentation/thompson>.
- Toulmin, S.E., 2003. The Uses of Argument. 2 ed., Cambridge University Press, Cambridge. URL: <https://www.cambridge.org/core/books/uses-of-argument/26CF801BC12004587B66778297D5567C>, doi:10.1017/CB09780511840005.
- TRUSTZONE A/S, 2024. Qualified Certificates - PSD2, QWACs and QSEALs. URL: <https://trustzone.com/qualified-certificates-psd2-qwacs-and-qseals/>.
- Tyndall, J., 2010. AACODS checklist URL: <https://scholar.google.com/scholar?cluster=15804274312996952804&hl=en&oi=scholar>.
- Vila-Henninger, L., Dupuy, C., Van Ingelgom, V., Caprioli, M., Teuber, F., Pennetreau, D., Bussi, M., Le Gall, C., 2024. Abductive Coding: Theory Building and Qualitative (Re)Analysis. *Sociological Methods & Research* 53, 968–1001. URL: <https://journals.sagepub.com/doi/10.1177/00491241211067508>, doi:10.1177/00491241211067508.
- Wazan, A.S., Laborde, R., Benzekri, A., Taj, I., 2024. Article 45 of the eIDAS Directive Unveils the need to implement the X.509 4-cornered trust model for the WebPKI, in: Proceedings of the 19th International Conference on Availability, Reliability and Security, ACM. doi:10.1145/3664476.3670900.
- Weigl, L., Amard, A., Codagnone, C., Fridgen, G., 2022. The EU's Digital Identity Policy: Tracing Policy Punctuations, in: Proceedings of the

15th International Conference on Theory and Practice of Electronic Governance, pp. 74–81. doi:10.1145/3560107.3560121.

Wilson, K., 2015. Distrusting New CNNIC Certificates. URL: <https://blog.mozilla.org/security/2015/04/02/distrusting-new-cnnic-certificates>.

PREPRINT

Appendix A. Additional Resources

Appendix A.1. Unfolding timeline of events

The controversy surrounding QWACs has unfolded through successive stages, each marked by shifts in how stakeholders framed their arguments. Broadly, this development can be described in three phases.

Phase 1: Before the 2021 amendment proposal. QWACs have existed in law since 2014, but their enforcement remained with private root programs. During this period, transport controls matured, and QWAC uptake was only marginal relative to DV, OV, and EV practices (Rescorla, 2018; Laurie et al., 2021; Amann et al., 2017). The CA/B BR continued to be the de facto reference for certificate issuance and RP behavior (CA/B Forum, 2025).

Phase 2: Proposal and talks (2021 to pre-leak). The eIDAS amendment proposal (Article 45) divided stakeholders. Browsers and civil society raised security and independence concerns, whereas QTSPs and governments emphasized verified identity visibility (Mozilla, 2021; Internet Society et al., 2023; ESD, 2022b). ETSI profiles outlined consumption paths in TLS certificates, parallel display, or out-of-band assertions for shifting debate toward governance, usability, and privacy-preserving identity checks (ETSI, 2025).

Phase 3: Leak and escalation (late 2023). A leaked draft intensified disputes. Critics warned of government pressure on trust stores, whereas supporters framed obligations as consumer protection (Muffett, 2023b; Mozilla et al., 2023). eIDAS 2.0 (2024/1183) now requires displaying attested identity without adding trust anchors or compromising TLS. Stakeholders view this as a pragmatic convergence, preserving transport security while enabling auditable, reliable, qualified identity (ETSI, 2023b).

Settlement and implementation (2024 onward). The eIDAS 2.0 Regulation (2024/1183) requires that, where a QWAC is available, user agents make attested identity visible to end users while avoiding interference with security protections and without mandating new trust anchors. ETSI guidance now documents decoupled consumption paths consistent with this direction, reducing friction with CAB Forum requirements (ETSI, 2025, 2023d; CA/B Forum, 2025). Browser statements characterized the outcome as averting earlier security risks, while supporters stressed that qualified identity gained a standardized presentation path (Mozilla, 2024; ESD, 2023c). Interviewees largely interpreted this as a pragmatic convergence that preserves transport security while testing whether a qualified identity can be made reliably retrievable and auditable in operator workflows.



Figure A.1: Timeline of QWAC related milestones and controversies.

Appendix A.2. Insights from interviews

[ST-S1] Strengthens Website Identity Assurance

- E10 - "I would like to know that I'm actually speaking to the Authentic Source, and the Authentic Source should also know that a QTSP is currently making the request. I have 27 member states, and I don't really want to do archaeological detective work to figure out whether the Authentic Source in Lithuania is one I can trust or whether I have some kind of problem here, because as a QTSP, I'm naturally gathering this information. I basically aggregate information and compile it into a certificate. However, you have to be accountable for each individual piece of information and know that it truly comes from a trustworthy source. That's the added value I provide, from my perspective, and we'd like to continue using QWACs in this area."

[ST-S2] Display Transparent Trust Indicator

- E8 - "I've not seen a study where I can say it's really a large scale, unbiased study of what people understand from the current indicators. And as I said, when I'm in current indicators, we need to be precise because Chrome has changed three or four times within two years. If you set and you let people understand with it, understand what you put you give a bit of time, then you do a study. Then we can see if it's well understood or not. And then we can actually try to improve it. But I'm not sure this effort has been done by browser trying to improve the visual indicators."

[ST-S3] Protect Users from Fraudulent Websites

- E15 - "We're already seeing the increase in fraud and victims, trapping into AI tricks because it's so easy to set up a fake website. Now I can go online and save a copy of a website and within seconds I have a copy that is perfect where there are no spelling mistakes where everything makes sense. And the only thing is needed domain name that looks like the other domain. Well, it's probably easy as well. So what we probably also want is we have these two significant approach. One is on TLS which is still under the browser rules and things like that. But you need to adhere to programs for compliance and things like that. It's QWAC."

[ST-C1] Do Not Enhance User Risk Awareness

- E5 - "Does it really work? And I think even if they succeed in doing this, will users really check these kind of signals? ... I think there was an honest effort to try it and it kind of failed. And then Europe just said, okay, we're going to make it mandatory. It has to do with what is the mental model of the user and how much effort does the users do?"

[ST-C2] Do Not Deliver Effective Trust Signals

- E8 - "All the CAs that were issuing QWAC were both QTSPs under EIDAS and both CAs recognised by the browsers. The only issue was that the source of trust for a qualified certificate by law is the European trusted list and the browsers say no source of the trust anchor for TLS has to be our own trust store. So there was a fundamental disagreement about who decides what can be trusted or not trusted."

[ST-C3] Degrade User Data Protection

- E6 - "whether it's surveillance or whether we call it privacy, there's a only a small difference for the user, right? Because in either case my private data, my private communication is exposed to some third party that I don't want it to be exposed to, whether that's a company or whether that's a state. "

[ST-C4] Introduce New Attack Vectors

- E15 - "We still need to see if that will get any adoption because it's way too complex for the average user to get that configured. The automation isn't prepared for that. We don't see that probably moving forward or anyone really using it."
- E13 - "I don't know from a user interface perspective if this is clear and clean enough because it feels to me that it's more complicated for the user to understand than the easy green bar."

[GA-S1] Promote Fair Competition in Digital Markets

- E15 - "you can automatically issue that certificate which means you will get competition because now the certificates can be issued at a much lower fee because there's no human work involved anymore. You can get it almost completely automated and you still need to go through that process, but that should be peanuts to walk through."

[GA-S2] Strengthen EU Digital Sovereignty

- E8 - "the digital sovereignty will not be much better if we go for the 2-QWAC approach."
- E1 - "We are Europe and it's not the intention that authorities or companies from outside Europe determine our legislative and technical framework. It's a kind of data sovereignty."

[GA-S3] Advance the EU Digital Single Market

- E9 - "They would say that of course they prefer there be a uniform level so that we do not have to build different things in seven regions."
- E3 - "To deploy apps like these widely, you must go through Google/Apple's tightly controlled App Store approval. While I'm glad they allowed certain contact tracing apps, it's unsettling how much power they hold over what gets deployed. This is partly why the EU's DMA now requires alternative app stores on EU devices."

[GA-C1] Facilitate Government Surveillance

- E1 - "To be honest, why QTSP is an independent third party. If the QTSP cheats, they are completely liable. So, there is no interest for the QTSP because the QTSP would question Its own existence, if they do surveillance us, because it will be checked. The QTSP is independent. This is a difference to Google. Google has an interest, a market interest. So surveil the people to analyse the behaviour of the people, just to sell them more product just to give them more advertisement, etcetera. So, there is a commercial interest of Google in doing this."
- E13 - "when I access my Gmail accounts. I probably do not expect that government agency certifies that I'm really talking to Google and not someone else. And even though, as I said, I'm generally believe that the governments in our Union are quite solid and democratic and are here in on our best interests, I think as a matter of principle, it's a good idea not to extend that trust to other areas."

[GA-C2] Undermine Neutral Global Trust Models

- E7 - "if you then create regulation that undermines the openness, the Google's and Meta's and so on, they used to deal with governments that come with special requests. I think it's more costly because they don't have all this legal infrastructure to deal with all these kind of things."

[GA-C3] Conflict with Existing Root Store Standards

- E1 - "One of the first versions of the thing in eIDAS was that the browsers were forced to accept the QWACs to put them in their root stores and without any possibility to control, to maintain, to avoid QWACs."

[CI-S1] Integrate Website AuthN into EU Trust Schemes

- E1 - "This is mainly what also did the QWAC solve, because we have independent legislator, we have an independent Certification of supervisory of the issue of qualified website authentication certificates and if you have an issue you can go to an independent court. So, we change the trust model from complete dependency on one team To a provable trust with the QWAC."

[CI-S2] Integrate with EU Cybersecurity Directives

- E10 - "I thought it made sense if it was organized at a European level in the sense that there would then be a decision made by ENISA or other responsible authorities and not by a private company. And when in doubt, the private company will do what someone tells it to anyway, which I think is an important point, especially in light of today's geopolitical developments"

[CI-S3] Strengthen Accountability and Transparency

- E13 - "When we establish a TLS communication, we don't usually receive multiple certificates. So the accountability is always with respect to the certificate authority that issues the certificate that we are using to make an encrypted change."
- E10 - "Accountability, exactly. And that doesn't exist, and now I have the alternative of going for the EV certificates. It's great that encryption is now taking place here, brilliant. It's brilliant, period."

But at the same time, we have the situation that all criminal websites are now computer-encrypted. So what have I achieved? It's simply that the transaction cannot be viewed by third parties. And I ask myself, how can we better protect the user? And I would say that the QWAC certificate is certainly a good alternative. From our point of view, QWACs can also be used in other areas, i.e., when it comes to communication between services, for example."

[CI-C1] Create Fragmented Trust Ecosystems

- E12 - "If you are a TLS only CA doing in the browsers, you have a web trust for CAs audit that is specifically for the TLS baseline requirements. Very specific criteria. And that is the entirety of that audit. If you're an ETSI CA, getting your qualified Audit. The regime incorporates how many different ETSI standards. 411-1, 411-2, 412-1 through 5-319-401. I mean, it just goes on and on and on. And so the number. And the audits are complicated, honestly. And if I'm a CA who is issued 2,000,000 qualified electronic signatures in a certificate and 100 QWACs, where's the audit emphasis go?"

[CI-C2] Undermine Technological Neutrality Principles

- E12 - "There's also some real difficulties involved in it, because we can invent this new idea. But the technology to make it work does not exist today."

[CI-C3] Increase Complexity and Costs for Website Operators

- E1 - "Yeah, there could be additional costs, but this is a question of the business model From The suppliers providing the software for those websites because if you have a one-person shop, typically You use templates, you use a provider providing templates for this website where you create a new website. You get your URL, and you publish. And in this case, it would be in some additional costs You may have to pay to this provider in this case, but to achieve European sovereignty and security, that should be done."

[CI-C4] Underperform Compared to Existing Measures

- E8 - "I mean certificate transparency In itself, I think it's not enough. You also need a mechanism to register to watch what is happening. Otherwise it's just a big website with a lot of information, if nobody's looking at it, you will not find the issue."

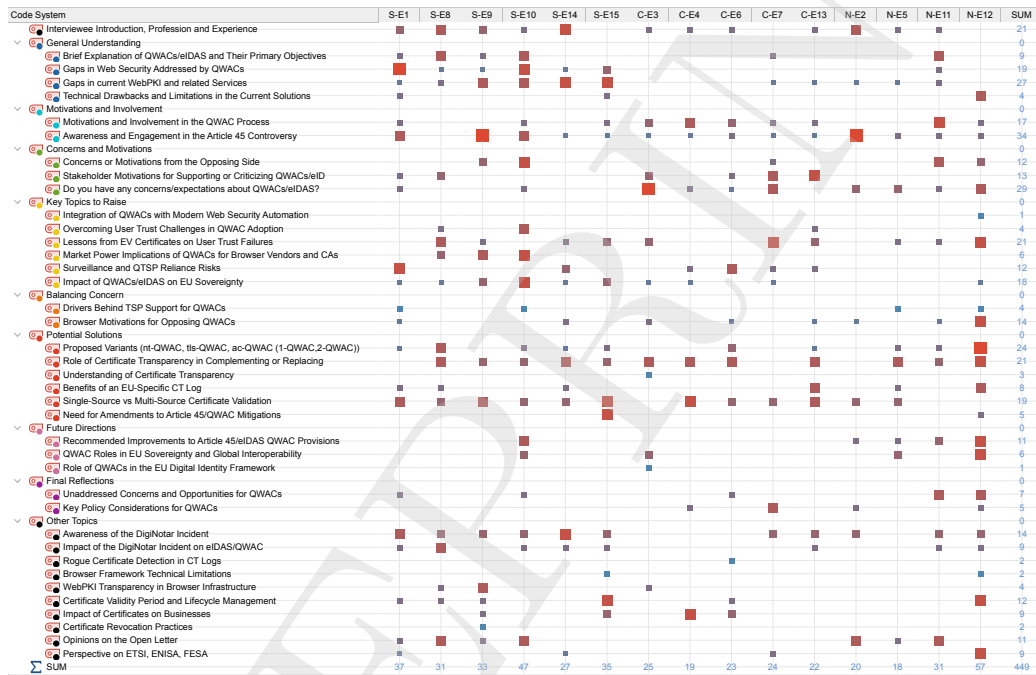


Figure A.2: Code Matrix Browser based on interview guide code system showing which topics were discussed by each expert with respective stance (S/C/N); symbol sizes were normalized per row in order to compare the relative prominence of arguments within individual expert interviews while controlling for differences in interview length

Code System	S-E1	S-E8	S-E9	S-E10	S-E14	S-E15	C-E3	C-E4	C-E6	C-E7	C-E13	N-E2	N-E5	N-E11	N-E12	SUM
ST																0
C4- Introduce New Attack Vectors																16
C3- Degrade User Data Protection																14
C2- Do Not Deliver Effective Trust Signals																24
C1- Do Not Enhance User Risk Awareness																16
S3- Protect Users from Fraudulent Websites																22
S2- Display Transparent Trust Indicator																27
S1- Strengthens Website Identity Assurance																28
GA																0
C3- Conflict with Existing Root Store Standards																17
C2- Undermine Neutral Global Trust Models																21
C1- Facilitate Government Surveillance																24
S3- Advance the EU Digital Single Market																1
S2- Strengthen EU Digital Sovereignty																30
S1- Promote Fair Competition in Digital Markets																18
CI																0
C4- Underperform Compared to Existing Measures																14
C3- Increase Complexity and Costs for Website Operators																18
C2- Undermine Technological Neutrality Principles																12
C1- Create Fragmented Trust Ecosystems																12
S3- Strengthen Accountability and Transparency																29
S2- Integrate with EU Cybersecurity Directives																10
S1- Integrate Website AuthN into EU Trust Schemes																21
SUM	18	13	21	37	37	30	27	23	29	22	18	6	28	35	30	374

Figure A.3: Interview Code Matrix of each interviewee for the argument-based code system

Appendix B. Argument Catalogue

This appendix provides the consolidated QWAC arguments derived from the multivocal corpus and expert interviews described in Section 3.3. For each argument, we document the Toulmin structure (Toulmin, 2003) by summarizing the *claim* and compiling the associated *grounds* (evidence referenced in sources), *warrants* (reasoning linking grounds to the claim), *backing* (supporting authority or contextual justification), *qualifiers* (scope conditions), and *rebuttals* (counter-considerations) as they appear across sources and stakeholder positions.

We organize the arguments into three heuristic and partially overlapping classes – *Security & Trust* [ST], *Governance & Authority* [GA], and *Compliance & Interoperability* [CI]. Each argument is assigned an identifier that encodes its class, stance (Supporters [S] or Critics [C]), and an incrementing number (e.g., [GA-C2]). To ensure coherence across arguments, all claims are interpreted relative to the shared premises defined in Section 4: *Legal Foundation* [P1], *Technical Specification* [P2], *Institutional Governance* [P3], and *Operational Enforcement* [P4].

Readers can use the Toulmin components documented for each argument to form their own judgement about how strongly a claim is supported in the corpus and how strongly it is contested. Concretely, assess whether the *grounds* are specific and sufficient, whether the *warrant* plausibly links those grounds to the claim, whether the *backing* is independent and relevant, whether *qualifiers* make scope and conditions explicit, and whether *rebuttals* materially weaken the core mechanism. If a quantitative evaluation is useful, apply a three-level rating: *Weak* if support is speculative or strongly contested, relies mainly on partisan backing, or is substantially undermined by rebuttals; *Substantial* if the mechanism is well articulated and supported but remains conditional on unsettled implementation/governance/enforcement choices and/or lacks outcome-level evidence; and *Strong* if grounds, warrants, and backing are mutually consistent and independently corroborated across stances, with explicit qualifiers and rebuttals that do not negate the core claim. We do not report author-assigned ratings.

For consistency, each argument summary follows the same Toulmin-oriented sentence structure: *Claim / Conclusion* has *Data / Evidence* since *Warrant / Explanation* on account of *Backing / Authority* presumably *Qualifier / Relativisation* unless *Rebuttal / Counter*, as illustrated in Figure 3.

[ST-S1] QWACs strengthen Website Identity Assurance

Claim/Conclusion: QWACs enhance website identity assurance by binding a vetted legal entity to the TLS server certificate used for website authentication, under public-law supervision and conformity assessment, which improves attribution and ex post accountability [P3], while leaving the underlying cryptography that ensures TLS unchanged [P2].

Data/Evidence: Only QTSPs listed in the EU LoTL may issue QWACs. QTSPs are subject to supervised policy and periodic conformity assessment against EN 319 401, EN 319 411-1 and-2, with additional controls for auditing bodies in TS 119 403-2. EN 319 412-4 defines the certificate profile for website certificates and, together with EN 319 412-3 and EN 319 412-1, harmonizes legal-entity identity semantics encoded in certificates. ETSI TS 119 411-5 (v2.1.1) specifies qualified issuance and two consumption patterns, 1-QWAC (in-chain) and 2-QWAC (binding), which keep the attested legal-entity data bound to the HTTPS endpoint. The Commission frames qualified status as adding supervised, enforceable organizational assurance beyond DV practice [P3] (European Commission, 2021d), and ENISA provides usage guidance and benefits caveats for QWACs (ENISA, 2016, 2017).

Warrant/Explanation: If organisational identity is verified by a supervised provider and cryptographically tied to the endpoint, relying parties can attribute interactions to a specific legal person and pursue remedies for mis-issuance or misuse [P3] (European Commission, 2021d; ENISA, 2017). Because QWACs do not modify the TLS protocol, any uplift is attributional and governance-based rather than cryptographic [P2] (ETSI, 2025; Rescorla, 2018).

Backing/Authority: Regulation (EU) 2024/1183, read with the eIDAS framework, establishes recognition, supervision, and display obligations relevant to qualified website authentication. EN 319 401 and EN 319 411-1/-2 set baseline assurance and policy requirements for QTSPs (ETSI, 2024a, 2023e). EN 319 412-4/-3/-1 standardise website certificate and legal-entity identity profiles (ETSI, 2023b,a,c). TS 119 411-5 v2.1.1 defines validation steps and relying-party consumption outcomes for 1-QWAC and 2-QWAC (ETSI, 2025). ENISA reports characterize the

qualified status as adding supervised organizational assurance and provide guidance on appropriate use (ENISA, 2016, 2017).

Qualifier/Relativisation: The assurance gain depends on audit quality and supervisory coordination, robust status and revocation handling, and faithful implementation of the specified profiles and consumption patterns [P1] (ETSI, 2024a, 2023g, 2025). The corpus documents these preconditions but does not include outcome studies that quantify mis-issuance reduction or phishing prevention specifically attributable to Article 45 QWAC implementations [P3] (European Commission, 2021d).

Rebuttals/Counter: Critics argue that QWACs do not strengthen TLS confidentiality or integrity and that 2-QWAC side-bindings add verification complexity relative to DV plus CT and short lifetimes [P4] (Rescorla, 2022; CCADB, 2020b; Mozilla, 2021, 2022; Certificate Transparency, 2020; Laurie et al., 2021). They further noted that prior EV-style identity indicators showed limited behavioral security benefits, motivating their demotion in browsers (Felt et al., 2016; Mozilla, 2021).

[ST-S2] **QWACs display Transparent Trust Indicator**

Claim/Conclusion: QWACs enable browsers to present a legally verified website-operator identity in a user-friendly manner [P1, P2, P3], because Article 45 of Regulation (EU) 2024/1183 mandates recognition and display of qualified identities [P1], and the underlying ETSI profiles (EN 319 412-4 and TS 119 411-5 v2.1.1) define standardized legal-entity attributes and consumption models that make such identity information interpretable and auditable across Member States [P2, P3].

Data/Evidence: Regulation (EU) 2024/1183 Art. 45 obliges user agents to recognize QWACs and to display the attested identity in a user-friendly manner [P1]. EN 319 412-4 and EN 319 412-3 define uniform identity attributes for website certificates and legal entities, including legal name, jurisdiction, and registration identifiers [P2]. TS 119 411-5 v2.1.1 specifies the 1-QWAC and 2-QWAC consumption outcomes and which identity semantics clients can retrieve for presentation [P2]. Root-store policies describe the operational context in which such identity signals are enforced in browsers [P4]. (Mozilla, 2023c)

Warrant/Explanation: When a binding legal duty to surface identity exists [P1] and the identity semantics and consumption models are standardized [P2], user agents can implement consistent and auditable identity signals across vendors and MS, supporting human comprehension and machine checks. However, prior usable-security research shows limited direct behavioral effect from positive identity indicators, so any transparency gains are likeliest in trained or policy-driven contexts [P3]. (Thompson et al., 2019; Felt et al., 2016)

Backing/Authority: The binding legal text and ETSI standards provide the authoritative basis for recognition, display duties, and uniform identity semantics [P1, P2]. ENISA guidance frames qualified identity disclosure as improving accountability and transparency [P3]. (ENISA, 2017) Supporter and academic analyses argue that harmonized semantics and program obligations can enable clearer four-corner trust and feasible UI treatments [P3, P4]. (Wazan et al., 2024; Entrust Corporation, 2021)

Qualifier/Relativisation: Neither the Regulation nor TS 119 411-5 prescribes a specific icon, placement, or wording; UX is left to implementers and program policies, so cross-vendor convergence depends on vendor choices and governance [P2, P3, P4]. (ETSI, 2025; Mozilla, 2023c) There are no controlled field studies in the corpus demonstrating improved recognition error rates from mandated QWAC identity display; prior EV results indicate low salience and limited behavioral change in general audiences [P3]. (Thompson et al., 2019; Felt et al., 2016)

Rebuttals/Counter: Critics argue that indicator-centric designs have yielded little user benefit, that additional verification paths such as 2-QWAC add complexity, and that UX divergence across browsers could blunt standardization benefits [P4]. (Mozilla, 2020; CCADB, 2020b; Helme, 2023; Thompson et al., 2019) Supporters counter that a mandated, harmonized signal plus machine-readable identity can still improve attribution, enterprise policy checks, and regulator oversight even if general-purpose positive indicators remain low salience [P1, P2]. (ENISA, 2017; Wazan et al., 2024; Entrust Corporation, 2021)

[ST-S3] **QWACs protect Users from Fraudulent Websites**

Claim/Conclusion: QWACs can reduce impersonation risk by providing vetted, machine-readable operator legal-entity identity at connection time,

which enables detection of brand-entity mismatches and supports attribution and redress [P1, P2].

Data/Evidence: Under the amended eIDAS framework, QWACs are recognized as qualified trust services and must be accepted in the EU [P1]. ETSI-conform issuance and auditing provide structured identity semantics and regular oversight for trust service providers [P2, P3] (ETSI, 2023g). Independent guidance describes QWAC identity attributes as inputs to user and tool decisions, including the organization name and registration data [P2] (ENISA, 2017, 2016). Two consumption patterns for QWACs are documented and widely discussed by root programs and implementers, enabling the retrieval of legal-entity data either in-chain or out-of-band so that relying parties and automated defenses can compare asserted brands with vetted operators [P2, P4] (CCADB, 2020b; Mozilla, 2023c). Empirical work shows that phishing sites typically do not replicate issuer and subject details, which creates detectable mismatches when a reliable identity is present [P2] (Drury and Meyer, 2019).

Warrant/Explanation: Because common phishing relies on look-alike domains with DV-only issuance, exposing QTSP-vetted legal-entity data at the moment of connection allows policy engines, security controls, or trained users to flag mismatches between claimed brands and verified operators, and to route incidents for accountability [P2, P3] (ENISA, 2017; Drury and Meyer, 2019). This identity signal complements, rather than replaces, the DV+CT baseline by adding an attribution layer that can be consumed by enterprise tooling and reputation systems [P4] (Certificate Transparency, 2020; Mozilla, 2023c).

Backing/Authority: The legal basis and obligations come from Regulation (EU) 2024/1183 and the original eIDAS 2014 text [P1]. Technical and conformity requirements are set by ETSI and related supervisory processes, while root-store programs define operational enforcement and incident handling [P2, P3, P4] (ETSI, 2023g; Mozilla, 2023c). ENISA documents the intended use of the QWAC identity in anti-fraud workflows and user guidance [P2] (ENISA, 2017, 2016).

Qualifier/Relativisation: The magnitude of phishing reduction depends on accurate brand-to-entity mapping, automation for discovery and re-

vocation, and trained consumption by relying parties; the corpus lacks controlled, outcome-level studies attributing incident reductions specifically to QWAC-based checks [P2, P4] (Tehrani et al., 2024; Felt et al., 2016; Thompson et al., 2019).

Rebuttals/Counter: Critics note that QWACs do not change TLS cryptography and that identity indicators have shown limited behavioral effect, arguing that DV+CT with short lifetimes is a simpler and already effective baseline [P4] (Rescorla, 2022; Certificate Transparency, 2020; Felt et al., 2016; Thompson et al., 2019; Mozilla, 2021). Interoperability notes also highlight the additional plumbing and operational complexity for 1-QWAC and 2-QWAC compared with domain-bound issuance semantics [P2, P4] (CCADB, 2020b).

[ST-C1] **QWACs do not enhance User Risk Awareness**

Claim/Conclusion: EV/QWAC-style identity indicators have little effect on user security behaviour and therefore provide limited incremental protection against phishing relative to a DV+CT baseline. [P2, P3, P4]

Data/Evidence: Users rarely notice or correctly interpret certificate-borne identity cues in realistic tasks, limiting behavioural leverage (Felt et al., 2016; Biddle et al., 2009). [P4] A large field experiment found that removing the EV indicator did not significantly change user behaviors tied to security outcomes, indicating the minimal influence of positive identity UI (Thompson et al., 2019). [P4] Major browser programs de-emphasized and then removed EV address-bar indicators without documented security regressions in program or practitioner materials (Feisty Duck, 2019; Hunt, 2019). [P3, P4] QWACs do not modify TLS transport guarantees; ETSI specifications define 1-QWAC and 2-QWAC models that add identity presentation or out-of-chain bindings rather than stronger channel security (ETSI, 2025; Rescorla, 2022). [P2] Empirical analyses report that certificate contents alone cannot reliably distinguish phishing from benign sites, while look-alike domains and social engineering dominate phishing success (Drury and Meyer, 2019). [P4] At Internet scale, the most effective programmatic controls against mis-issuance and impersonation are DV under BR, comprehensive CT logging and monitoring, and short lifetimes with automated renewal (Laurie, 2014; Laurie et al., 2021; Aas et al., 2019; Mozilla, 2023c; Certificate Transparency, 2020). [P3]

Warrant/Explanation: If end users and many relying parties do not reliably perceive or use legal-entity signals at decision time, then strengthening organisational vetting or mandating display will not materially change click-through, form-submission, or trust behaviours that determine phishing exposure; under those conditions, EV/QWAC indicators add little beyond domain control and existing browser protections (Felt et al., 2016; Thompson et al., 2019; Mozilla, 2021). [P3, P4]

Backing/Authority: Program decisions and practitioner analyses document EV UI deprecation and the rationale for deprioritising positive identity indicators (Feisty Duck, 2019; Hunt, 2019) [P3]. Technical analyses explain that QWACs leave TLS semantics unchanged and primarily introduce additional identity-display or binding logic (Rescorla, 2022). The ETSI EN and TS profiles specify identity semantics and 1-QWAC/2-QWAC consumption models but do not strengthen channel cryptography [P2]. WebPKI controls, such as CT and short lifetimes, are widely deployed and enforced by root store policy, providing transparent issuance oversight and reducing key-compromise dwell time (Laurie et al., 2021; Certificate Transparency, 2020; Aas et al., 2019; Mozilla, 2023c) [P3].

Qualifier/Relativisation: Behavioral impact may be larger in managed or high-assurance contexts where users are trained to verify legal-entity attributes, where enterprise policy engines consume identity fields, or where UI patterns converge; however, we found no controlled studies demonstrating population-scale phishing-rate reductions specifically attributable to mandated QWAC display (Mozilla, 2021; Rescorla, 2022). [P1, P2, P4]

Rebuttals/Counter: Supporters argue that Article 45 recognizes QWACs and requires a user-friendly display, and that harmonized ETSI profiles could enable training, automated checks, and clearer accountability that deter impersonation (ETSI, 2023d, 2025; Entrust Corporation, 2021; Entschew et al., 2022). [P1, P2]

[ST-C2] **QWACs do not deliver Effective Trust Signals**

Claim/Conclusion: EV/QWAC-style positive identity indicators are rarely noticed and are poorly understood by end users; controlled and field

measurements have not shown these indicators to improve real user security decisions; after major browsers demoted or removed EV address-bar UI, studies observed no measurable degradation in user outcomes. Therefore, indicator-centric protections add little beyond DV plus existing CT-based accountability and browser safeguards [P2, P4].

Data/Evidence: In 2019, Chrome and Firefox announced and shipped removal of prominent EV UI in the address bar (Feisty Duck, 2019; Hunt, 2019). A large-scale field experiment found that removing EV indicators produced minimal changes in user behavior and did not reduce security (Thompson et al., 2019). Multiple user-studies show that positive identity cues are seldom noticed and are often misinterpreted (Felt et al., 2016; Stojmenovic and Biddle, 2018; Biddle et al., 2009). QWACs do not alter TLS cryptographic semantics, so any benefit must stem from consumption and presentation of identity data rather than stronger transport guarantees (Rescorla, 2022) [P2]. Meanwhile, Certificate Transparency and short certificate lifetimes have already improved detection, accountability, and incident response without relying on the user interpretation of positive indicators (Laurie et al., 2021; Certificate Transparency, 2020) [P4].

Warrant/Explanation: If users and relying parties neither reliably notice nor correctly act on certificate-borne identity signals at the decision point, then mandating or enhancing such signals (EV or QWAC) will not meaningfully change the click-through or submission behaviors that drive phishing success (Felt et al., 2016; Thompson et al., 2019). Because QWAC does not change the TLS security properties of the connection, its potential benefit depends on consistent UI consumption and reliable relying-party logic rather than stronger cryptography (Rescorla, 2022) [P2]. Under these conditions, the marginal protection over DV plus CT and existing browser mitigations is limited (Laurie et al., 2021; Certificate Transparency, 2020) [P4].

Backing/Authority: Program communications and practitioner reports document EV UI de-emphasis and removal (Hunt, 2019; Feisty Duck, 2019). Usable-security research consistently finds weak notice and comprehension for positive identity indicators (Felt et al., 2016; Stojmenovic and Biddle, 2018; Biddle et al., 2009). Empirical analyses show widespread phishing over HTTPS, undermining naive identity-salience heuristics

(Drury and Meyer, 2019). Technical analyses argue that QWAC leaves TLS semantics unchanged and shifts any gains to identity consumption logic and UI (Rescorla, 2022). Standard baselines and operational controls already strengthen accountability via CT and related mechanisms (Laurie et al., 2021; Certificate Transparency, 2020) [P2, P4].

Qualifier/Relativisation: Effectiveness is context-dependent. In managed environments with training, policy engines, and automation that consume harmonized identity attributes, salience and utility could increase (ETSI, 2023d,f) [P2, P3]. Article 45 obligations to recognize a qualified identity and present it in a user-friendly manner may enable a more consistent UI and education in the EU context (ESD, 2023a) [P1, P2]. However, we found no peer-reviewed, population-scale evidence that mandated QWAC display reduces phishing relative to the current DV+CT baselines [P4].

Rebuttals/Counter: Supporters contend that, unlike voluntary EV, Article 45 creates a legal duty for recognition and user-friendly surfacing of harmonised identity attributes, which can support consistent presentation, user education, and automated checks (ESD, 2023a) [P1]. They further point to ETSI profiles that standardize identity semantics and guidance for coexistence with browser controls, arguing that this enables reliable consumption pathways across UAs and enterprises (ETSI, 2023d,f) [P2, P3].

[ST-C3] **QWACs degrade User Data Protection**

Claim/Conclusion: QWAC validation and display workflows can increase exposure of user browsing metadata by introducing additional parties (for example, QTSPs and supervisory infrastructure) into validation and status-checking paths, potentially revealing which sites a user visits and when [P1, P3]. Relative to DV+CT, baseline auditing relies on CT logs and verifiers that do not require per-visit online lookups.

Data/Evidence: ETSI TS 119 411-5 specifies consumption models, including a side-binding in which a QWAC is validated alongside a conventional WebPKI cert, a pattern that can require additional binding discovery and processing steps. (ETSI, 2025) CCADB documents note extra verification, plumbing, and interoperability concerns when identity is

bound outside the normal chain, implying additional discovery or status interactions. (CCADB, 2020a) Mozilla risk analyses describe flows that fetch side-bindings or consult issuer and registry endpoints, increasing the set of entities that can observe visit metadata [P1, P3]. (Mozilla, 2022, 2023a) Empirical work shows that legacy online status mechanisms (OCSP, CRL) leak per-site queries unless mitigated by stapling, demonstrating the privacy cost of live status checks. (Berbecaru and Lioy, 2023; Eastlake 3rd, 2011) Public briefings and practitioner analyses warn that state-aligned QTSPs could correlate issuance logs with access telemetry if per-visit queries occur [P3]. (Scientists and NGOs, 2023a; Helme, 2023; Muffett, 2023a) By contrast, CT is log-based and can be validated without per-visit lookups. (Laurie et al., 2021; Certificate Transparency, 2020)

Warrant/Explanation: If browsers perform extra online discovery or revocation checks for qualified-identity artifacts at page load, those requests reveal timing and destination context to contacted QTSPs or registries, enabling inference and cross-site correlation even under TLS [P1, P3]. (Eastlake 3rd, 2011; Berbecaru and Lioy, 2023) Since CT auditing can be validated without per-visit traffic, introducing qualified-identity plumbing plausibly enlarges the metadata surface relative to DV+CT. (Laurie et al., 2021; Certificate Transparency, 2020)

Backing/Authority: ETSI TS 119 411-5 defines the relevant consumption and validation models, including side-binding. (ETSI, 2025) CCADB highlights the added plumbing and binding challenges beyond standard chain validation. (CCADB, 2020a) Mozilla and NGO materials articulate the privacy and interception risks of mandated recognition and display under Article 45, explaining how added checks create new observation points [P1, P3]. (Mozilla, 2022, 2023a; Scientists and NGOs, 2023a) Practitioner commentary details operational patterns consistent with remote checks. (Helme, 2023)

Qualifier/Relativisation: The magnitude of exposure depends on the implementation: neither eIDAS 2.0 nor ETSI TS 119 411-5 mandates per-visit online validation, and offline validation is permissible [P1, P2]. Privacy-preserving status models such as OCSP stapling can reduce leakage relative to live OCSP queries, and local verification of qualified bindings can avoid network contact with QTSPs [P4]. (Eastlake 3rd,

2011; Berbecaru and Liroy, 2023) CT remains a log-based mechanism that can be validated without per-visit lookups. (Laurie et al., 2021)

Rebuttals/Counter: Supporters emphasize that Article 45 requires recognition and user-friendly display, not phone-home lookups, and that ETSI profiles support offline validation so that well-engineered implementations need not increase the DV+CT network footprint [P4]. Proponents further argue that broad privacy-exposure claims conflate optional implementations with required behavior. (ESD, 2022a,b, 2023c)

[ST-C4] **QWACs introduce new Attack Vectors**

Claim/Conclusion: QWACs risk degrading web security by adding validation paths and identity-UI dependencies that do not change TLS semantics, which can introduce new failure modes and a false sense of assurance relative to DV+CT baselines [P2, P4].

Data/Evidence: (i) TLS transport guarantees remain unchanged under QWACs, so any purported benefit must arise from identity presentation rather than cryptographic hardening (ETSI, 2025; Rescorla, 2022) [P2]. (ii) Both 1-QWAC and 2-QWAC require additional verification plumbing (for example, side-bindings, attribute extraction, discovery, and revocation or status checks), expanding the attack surface and operational complexity (ETSI, 2025; CCADB, 2020a) [P2, P4]. (iii) Prior reliance on positive identity UI (EV indicators) was demoted or removed by major browsers without observed safety improvements, and empirical work found limited user protection from such cues (Feisty Duck, 2019; Drury and Meyer, 2019; Felt et al., 2016) [P4]. (iv) Legal recognition and display duties can constrain fast incident responses or impose fragile UI obligations during active failures, creating inconsistent or exploitable states (Mozilla, 2021; Mozilla et al., 2023; Mozilla, 2023c; CA/B Forum, 2020) [P1, P3, P4].

Warrant/Explanation: If transport security is unchanged, added verification and presentation logic raises complexity and with it the probability of implementation bugs, misconfigurations, or downgrade paths; meanwhile, salient-but-misunderstood identity UI can induce over-trust, so indicators under failure can worsen effective security (Rescorla, 2022; CCADB, 2020a; Felt et al., 2016) [P2, P4].

Backing/Authority: Standards and program materials document that QWACs define consumption outcomes without altering TLS (ETSI, 2025; Rescorla, 2022) [P2]; CCADB discusses extra plumbing and interoperability risks for non-chain bindings (CCADB, 2020a) [P4]; EV UI was deliberately removed with rationale and without measured security benefit (Feisty Duck, 2019; Hunt, 2019; Drury and Meyer, 2019) [P4]; CT already provides transparent issuance oversight as a mature baseline (Certificate Transparency, 2020) [P4].

Qualifier/Relativisation: Risk magnitude is implementation-dependent: the final Regulation and ETSI profiles mandate recognition and define outcomes but do not require per-visit online lookups or prescribe brittle UI; careful engineering (offline validation, stapled status, CRLite-style revocation, consistent UX) can limit new failure modes (ETSI, 2025; Mozilla, 2023c) [P1, P2, P4].

Rebuttals/Counter: Supporters argue QWACs align with established baselines and supervisory controls, define controlled 1-QWAC and 2-QWAC consumption, and leave transport security unchanged; they claim that audits, harmonised semantics, and properly engineered flows can avoid brittle dependencies and UI hazards (ETSI, 2025, 2024a, 2023b; Entschew et al., 2022; Wazan et al., 2024; ENISA, 2017) [P1, P2, P3].

[GA-S1] **QWACs promote Fair Competition in Digital Markets**

Claim/Conclusion: Mandated browser recognition and user-friendly display of QWACs can promote fairer competition in digital trust markets by reducing reliance on non-EU browser root-store gatekeepers and enabling EU-supervised QTSPs to have their verified identities consistently consumed across the single market [P1, P3].

Data/Evidence: Article 45 of Regulation (EU) 2024/1183 requires user agents to recognise QWACs issued by QTSPs listed on Member State trusted lists aggregated in the EU List of Trusted Lists and to display the attested identity in a user-friendly manner [P2]. The Commission’s impact assessment and supporting study articulate competitiveness and internal market objectives for the revision [P3] (European Commission, 2021b,e). Supporter materials argue that mandatory recognition reduces the discretionary power of global root-store operators and levels

access to relying parties [P1, P3] (Bitkom, 2020; Bundesdruckerei, 2022; ESD, 2022b, 2023c; Schwalm, 2023; Identityum, 2022). Harmonised conformity and supervision frameworks apply to EU QTSPs (general policy EN 319 401; CA requirements EN 319 411-1; identity semantics EN 319 412-4 and -5) [P2] (ETSI, 2024a, 2023d,e). Evidence of active EU QTSPs issuance and market offerings for QWACs exists (e.g., certSIGN repository; vendor product pages) [P3] (certSIGN, 2024; Sectigo Limited, 2024; TRUSTZONE A/S, 2024).

Warrant/Explanation: If access to browser consumption of organisational identity is guaranteed by law for entities audited under a common EU regime and if identity surfacing by user agents is mandatory, then competitive access shifts away from discretionary inclusion by private browser root programs toward compliance with transparent, public-law criteria, which can reduce entry barriers and mitigate dependence on non-EU gatekeepers [P1, P2] (Mozilla, 2023c; CCADB, 2020b).

Backing/Authority: The legal basis defines recognition, display, and TL/LoTL governance [P2]. The Commission’s impact assessment provides a competitiveness rationale [P3] (European Commission, 2021b). The ETSI standards specify uniform provider requirements and identity data models that make the outputs substitutable for relying parties [P2] (ETSI, 2024a, 2023b,e). Supporter analyses describe expected level-playing-field effects [P3] (Bitkom, 2020; ESD, 2023c; Bundesdruckerei, 2022).

Qualifier/Relativisation: The competition effect is contingent on implementation and uptake: statutory recognition does not itself create demand for qualified identity, many sites may remain on DV/OV issuance, and the behavioural effect of identity UI has been limited in experiments [P4] (Thompson et al., 2019; Biddle et al., 2009; Helme, 2021). Interoperability with existing WebPKI practices and consistent supervisory application across Member States are prerequisites for any levelling effect [P2, P3] (Rescorla, 2022; CCADB, 2020b; Mozilla, 2021).

Rebuttals/Counter: Critics warn that statutory recognition can politicise inclusion and weaken independent browser security governance, or create a parallel path with extra costs that burden smaller providers, without materially changing browser market power [P1] (Mozilla, 2021, 2023b;

Mozilla et al., 2023; Rescorla, 2022; Internet Society et al., 2023). They further argued that certificate identity content is a weak anti-phishing signal and that past identity indicators had little measurable user effect [P4] (Drury and Meyer, 2019; Thompson et al., 2019).

[GA-S2] **QWACs strengthen EU Digital Sovereignty**

Claim/Conclusion: QWACs strengthen EU digital sovereignty by shifting recognition and presentation of website identity from private, extra-EU root-stores toward EU-supervised trust services, embedding organisational authentication within a public-law framework. [P1, P3, P4]

Data/Evidence: Regulation (EU) 2024/1183 obliges user agents to recognise QWACs issued by QTSPs on Member-State Trusted Lists aggregated in the EU List of Trusted Lists and to display the attested identity "in a user-friendly manner". [P1] The trusted list architecture originates in the eIDAS 2014 framework and remains the registry basis for the recognition of qualified trust services. [P1, P3] Continuous supervision and periodic conformity assessment are anchored in ETSI EN 319 401 and related audit schemes for CABs (ETSI, 2024a, 2023g). [P2, P3] ETSI TR 119 411-5 provides guidance for coexistence with browser trust stores and consumption outcomes aligned to Article 45 (ETSI, 2023f). [P2] Browser root programs illustrate the baseline of private, extra-EU discretion over recognition, distrust and incident handling (Mozilla, 2023c; Wilson, 2015). [P4] Supporter materials explicitly frame QWACs as an EU-centred trust layer that advances sovereignty (Bitkom, 2020, 2022; Bundesdruckerei, 2022; ESD, 2023c). [P3]

Warrant/Explanation: If law binds user agents to recognise and surface identities attested by EU-supervised QTSPs, then decision rights over what organisational identity is accepted and how it is presented move from private root-store owners to criteria defined and enforced under EU public law (ETSI, 2024a). [P1, P3] Because root programs have historically set and enforced recognition and distrust policies unilaterally (Mozilla, 2023c; Wilson, 2015), statutory recognition plus supervised issuance and audit reduce dependency on non-EU gatekeepers for policy change, incident criteria, and identity semantics (ETSI, 2023f). [P1, P2, P4] Analytical work on Article 45 situates this shift within a 4-cornered trust model that relocates validation authority into an EU-governed layer (Wazan et al., 2024). [P3]

Backing/Authority: Binding authority is provided by Regulation (EU) 2024/1183 for recognition and display obligations. [P1] Supervision, certificate content, and assurance baselines are specified in the ETSI EN 319 401 and EN 319 412 parts (ETSI, 2024a, 2023c). [P2] Coexistence and consumption guidance are documented in ETSI TR 119 411-5 (ETSI, 2023f). [P2] Governance analyses and stakeholder positions further substantiate the sovereignty framing (Entschew and van Brouwershaven, 2024; Wazan et al., 2024; Bitkom, 2020; ESD, 2023b). [P3]

Qualifier/Relativisation: The sovereignty gain is contingent on effective cross-border supervision and audit quality, interoperable implementation with global WebPKI operations, and consistent, secure UI behaviour (ETSI, 2024a, 2023f; ENISA, 2017). [P2, P3] Mandated recognition does not alter browser control over software delivery, update channels, and UI conventions, which can limit the magnitude of the practical shift (Mozilla, 2023c; Feisty Duck, 2019). [P4] Robust transparency and monitoring in the existing WebPKI remain relevant complements (Certificate Transparency, 2020). [P4]

Rebuttals/Counter: Critics warn that statutory recognition can politicise inclusion or expand surveillance leverage and that creating a parallel path may slow incident response, while root programs still retain decisive enforcement powers (Electronic Frontier Foundation (EFF), 2022; Mozilla, 2021; Claburn, 2023; Mozilla, 2023a,c). [P4]

[GA-S3] **QWACs advance the EU Digital Single Market**

Claim/Conclusion: QWACs provide a harmonised, cross-border framework for organisational website identity that Union law requires to be recognised by user agents and presented in a user-friendly manner, establishing a single pathway for identity display across Member States (ETSI, 2023b) [P1, P2, P3]. By reducing informational frictions in cross-border onboarding and discovery, this measure supports Digital Single Market integration and EU competitiveness [P1, P3].

Data/Evidence: Regulation (EU) 2024/1183 obliges application providers to recognise QWACs issued by QTSPs on Member-State Trusted Lists aggregated via the EU List of Trusted Lists, and to display the attested

identity in a user-friendly manner [P1, P3]. ETSI EN 319 401 sets supervised TSP requirements and lifecycle controls, while EN 319 412-4 harmonises website-certificate identity attributes (for example, organisationName, jurisdiction, and registration identifiers) that map to a legal entity (ETSI, 2024a, 2023b) [P2, P3]. TS 119 411-5 V2.1.1 specifies consumption outcomes for both 1-QWAC and 2-QWAC bindings, so relying parties can consistently retrieve and present identity data (ETSI, 2025) [P2, P3]. The Commission’s Impact Assessment frames the Article 45 revisions as instruments for internal market integration and international competitiveness (European Commission, 2021b) [P1]. Supporter materials explicitly link statutory recognition and a harmonized identity display to lower cross-border onboarding costs and improved discoverability of trustworthy EU services (Bitkom, 2020, 2022; Bundesdruckerei, 2022; ESD, 2023c) [P3]. The ENISA guidance describes the appropriate use of QWACs to promote consumer trust by verifying controllers (ENISA, 2017) [P2, P3].

Warrant/Explanation: If organisational website identity is verified under common EU criteria, surfaced consistently in browsers, and backed by continuous supervision, then cross-border provision faces fewer informational and trust asymmetries (who is the controller; under which legal identity), lowering search and compliance costs and facilitating DSM-wide provision and uptake of online services (European Commission, 2021b; ETSI, 2024a, 2025) [P1, P2, P3]. Harmonized identity semantics in EN 319 412-4 and supervised issuance under EN 319 401 mean that relying parties and regulators can attribute processing to a specific legal person uniformly across Member States (ETSI, 2023d, 2024a) [P2, P3].

Backing/Authority: The amended Regulation codifies recognition and display obligations and the MS-TL/LoTL governance model [P1, P3]. ETSI EN 319 401 and EN 319 412-4 provide technical and assurance controls, with deployment and consumption outcomes defined in TS 119 411-5 V2.1.1 (ETSI, 2024a, 2023b, 2025) [P2, P3]. The Commission’s Impact Assessment and policy communications articulate DSM and competitiveness objectives (European Commission, 2021b, 2022) [P1]. Advocacy materials from industry and practitioner groups provide additional rationales linking harmonised identity attestation to

internal-market benefits (Bitkom, 2020, 2022; Bundesdruckerei, 2022; ESD, 2023c) [P3].

Qualifier/Relativisation: Realising single-market benefits depends on uniform supervisory practice and interoperable implementations, especially for 2-QWAC bindings and associated retrieval pathways (ETSI, 2025) [P2, P3]. Practical convergence also requires a stable and comprehensible UI across browsers; however, root programs retain control over client security posture and user-interface decisions (Mozilla, 2023c; Feisty Duck, 2019) [P4]. Adoption by service providers is necessary, and the Regulation itself does not induce demand outside the EU (European Commission, 2022) [P1].

Rebuttals/Counter: Critics argue a parallel, mandated trust path could fragment the global WebPKI, conflict with root-store policies, and add compliance costs that offset DSM efficiency gains (Mozilla, 2023c; CCADB, 2020a; Feisty Duck, 2022; Helme, 2023) [P4]. They also warn that politicized inclusion or heterogeneous UI could undermine the predictability and identity salience that single-market integration seeks to achieve (Mozilla, 2023c; Feisty Duck, 2019) [P4].

[GA-C1] **QWACs facilitate Government Surveillance**

Claim/Conclusion: Mandatory recognition and display of QWACs can enable government-influenced trust anchors to facilitate targeted interception or broader interference with web traffic. [P1, P3, P4]

Data/Evidence: Article 45 of Regulation (EU) 2024/1183 obliges user agents to recognise QWACs and to display the attested identity in a user-friendly manner. [P1] QWACs are issued by QTSPs listed on Member-State Trusted Lists that are aggregated via the Union List of Trusted Lists, establishing a legally curated set of trust anchors. [P3] Article 45a sets a supervisory precautionary-measures process that can constrain immediate unilateral distrust by user agents while coordination with competent authorities occurs. [P1, P4] Browser and civil-society analyses describe a scenario in which a state-aligned or state-compelled QTSP could issue a certificate for a target domain that user agents initially accept until countermeasures are coordinated. (Mozilla, 2023a; Internet Society et al., 2023) [P1, P3, P4] Press reports and expert

commentary reiterate this risk and its implications for interception. (Claburn, 2023; Rescorla, 2022) [P4]

Warrant/Explanation: If inclusion on MS-TLs and the LoTL yields a recognised trust anchor that user agents must accept, and immediate unilateral distrust is procedurally mediated via Article 45a, then a mis-issued or compelled certificate can enable man-in-the-middle interception until revocation or distrust is effected. [P1, P3, P4] Certificate Transparency provides post-issuance detection and auditing but does not itself prevent the first successful acceptance of a presented certificate. (Laurie et al., 2021; Certificate Transparency, 2020) [P2, P4]

Backing/Authority: Legal obligations for recognition, display, and supervisory procedures are set by Regulation (EU) 2024/1183. [P1] The institutional trust-list framework and LoTL aggregation come from Regulation (EU) No 910/2014. [P3] General policy and qualified-certificate requirements for TSPs are specified by ETSI EN 319 401 and EN 319 411-1/-2. (ETSI, 2024a, 2023b,e) [P2] Browser coexistence guidance and consumption models are described in ETSI TR 119 411-5 and ETSI TS 119 411-5 v2.1.1. (ETSI, 2023f, 2025) [P2] Root-store governance and emergency distrust capabilities are governed by the Mozilla Root Store Policy. (Mozilla, 2023c) [P4] Independent critiques and reportage articulate the specific interception scenario under Article 45/45a. (Mozilla, 2023a; Internet Society et al., 2023; Rescorla, 2022; Claburn, 2023) [P4]

Qualifier/Relativisation: The risk is conditional: Article 45 does not authorise decryption or mandate changes to TLS, and browsers can continue to enforce CT logging, revocation, and emergency distrust policies. (Laurie et al., 2021; Mozilla, 2023c) [P1, P4] ETSI TS 119 411-5 describes a 2-QWAC binding model that preserves WebPKI transport semantics rather than introducing an alternative cryptographic channel. (ETSI, 2025) [P2] Real-world exploitation therefore depends on state capabilities, QTSP governance, and the latency of detection and coordinated distrust. [P1, P2, P4]

Rebuttals/Counter: Supporters argue that QWACs do not bypass TLS or browser security policy and that emergency distrust remains available;

they further claim that qualified status raises accountability under supervision and that ETSI profiles align with CA/B Forum Baseline Requirements. (Entschew et al., 2022; Martius et al., 2024; ETSI, 2023b,e; CA/B Forum, 2025) [P1, P2, P4]

[GA-C2] QWACs undermine Neutral Global Trust Models

Claim/Conclusion: QWAC obligations risk conflicting with established browser root-store policies and CA/Browser Forum baselines by legally mandating recognition and identity display under MS-TL/LoTL governance and by conditioning distrust on supervisory procedures rather than program discretion [P1, P3, P4].

Data/Evidence: Regulation (EU) 2024/1183 amends Article 45 to require user agents to recognise QWACs and display the attested identity in a user-friendly manner, and adds Article 45a on precautionary measures for distrust [P1]. Stakeholder analyses interpret Article 45(1b) as constraining additional mandatory requirements beyond Annex IV for user agents [P1, P3] (Mozilla, 2023b; Mozilla et al., 2023). Browser root programs state they enforce their own inclusion, incident-response, and policy criteria at their discretion, including BR compliance and CT expectations [P3, P4] (Mozilla, 2023c; CA/B Forum, 2025; Certificate Transparency, 2020). ETSI TS 119 411-5 v2.1.1 defines two consumption patterns for QWACs, including a 2-QWAC model that binds identity in a separate qualified artifact alongside a conventional TLS server certificate [P2] (ETSI, 2025). The CCADB and CA/B Forum memoranda document interoperability and verification gaps when organizational identity is bound out-of-chain relative to existing root-store hooks [P2, P4] (CCADB, 2020a; CA/B Forum, 2020).

Warrant/Explanation: When recognition and display are mandated and distrust is mediated through public-law trust lists and supervisory processes, while browser programs rely on private program policies, CT logging, and rapid unilateral enforcement tied to TLS-chain artefacts, mismatched authorities and artefacts predictably create governance friction and veto points that can conflict with root-store standards and incident-response practice [P1, P3, P4] (Mozilla, 2023c; Laurie et al., 2021; Certificate Transparency, 2020).

Backing/Authority: Mozilla policy papers and a joint civil-society research letter warn that Article 45 narrows user-agent discretion relative to established WebPKI governance [P1, P3] (Mozilla, 2021; Scientists and NGOs, 2023b). The Regulation itself encodes recognition, display, and supervisory roles that are distinct from root program practice [P1]. The CCADB and CA/B Forum notes substantiate the technical and policy divergence introduced by out-of-chain identity bindings under 2-QWAC [P2, P4] (CCADB, 2020a; CA/B Forum, 2020).

Qualifier/Relativisation: The degree of conflict is contingent: ETSI TS 119 411-5 aims to align qualified profiles with BR practice and preserves TLS semantics in the 2-QWAC model, so transport security and CT mechanisms remain unchanged [P2, P4] (ETSI, 2025; Laurie et al., 2021; Certificate Transparency, 2020). Article 45 is technology-neutral, and Article 45a does not categorically preclude urgent mitigations by user agents [P1].

Rebuttals/Counter: Proponents argue there is no material conflict because TS 119 411-5 defines browser-facing outcomes compatible with BR and leaves WebPKI transport semantics intact [P2, P4] (ETSI, 2025). Advocacy and industry statements claim that QWACs strengthen European sovereignty and consumer protection and can coexist with browser controls [P2, P3] (Bitkom, 2020; ESD, 2023c).

[GA-C3] **QWACs conflict with Existing Root Store Standards**

Claim/Conclusion: Mandatory recognition and display of QWACs risks shifting trust-anchor governance from browser root programs to statutory MS-TL/LoTL listings, enabling politicised inclusion that undermines browser independence [P1, P3, P4].

Data/Evidence: Article 45 of Regulation (EU) 2024/1183 obliges user agents to recognise QWACs and to display the attested identity in a user-friendly manner; Article 45a establishes supervisory procedures around distrust [P1]. Recognition is attached to QTSPs listed on Member-State Trusted Lists and aggregated via the Union List of Trusted Lists, that is, statutory listings distinct from browser root programs (European Commission, 2021a; ETSI, 2024a) [P3]. ETSI TS 119 411-5 defines two consumption models for qualified website authentication: 1-QWAC (server-cert path) and 2-QWAC (binding alongside a WebPKI

chain) (ETSI, 2025) [P2]. Browser and civil-society sources warn that statutory listing plus mandatory UX treatment can displace private root-store policy with political inclusion decisions (Mozilla, 2021; Scientists and NGOs, 2023a; Information Security Media Group, 2023; Recorded Future News, 2023; Mozilla, 2023a; Electronic Frontier Foundation, 2023) [P3, P4].

Warrant/Explanation: When acceptance and identity display are legally tied to MS-TL/LoTL listing rather than to vendor-defined root-store baselines, Member-State processes gain leverage over which QTSPs function as effective trust anchors for identity signalling and, under 1-QWAC, potentially for TLS termination (ETSI, 2025) [P1, P2, P3]. This reduces autonomy and rapid incident response discretion in browser root programs, which articulate independent inclusion, audit, and emergency distrust criteria (Mozilla, 2023c) [P4]. Consequently, statutory governance can pressure or preempt root-program judgment, shifting trust-anchor governance toward MS-TL/LoTL control (Mozilla, 2021; Scientists and NGOs, 2023a) [P3, P4].

Backing/Authority: Legal obligations for recognition, display, and supervision are set by Articles 45 and 45a [P1]. The accreditation and listing mechanics for QTSPs are defined in ETSI EN 319 401 (ETSI, 2024a) [P2, P3]. Consumption models that determine how qualified identity interacts with WebPKI transport are specified in ETSI TS 119 411-5 (ETSI, 2025) [P2]. Independent browser-governed baselines and emergency distrust powers are codified in the Mozilla Root Store Policy (Mozilla, 2023c) [P4]. Civil-society and industry reporting warn of politicisation risks and policy conflicts if statutory listings effectively override root-store policy (Scientists and NGOs, 2023a; Information Security Media Group, 2023; Recorded Future News, 2023; Electronic Frontier Foundation, 2023) [P3].

Qualifier/Relativisation: Displacement of root-store policy is not inevitable: the Regulation does not mandate modifying browser root-store contents or prescribe 1-QWAC over 2-QWAC [P1]. TS 119 411-5 provides a 2-QWAC binding that preserves WebPKI transport anchoring while surfacing the qualified identity, limiting the direct pressure on root programs [P2, P4]. Therefore, the risk magnitude depends on deployment

choices (1-QWAC vs. 2-QWAC), supervisory practice, and vendor implementation decisions (Mozilla, 2023c) [P2, P3, P4].

Rebuttals/Counter: Supporters argue that Article 45 recognises qualified identity without dictating root-store content [P1]; that browsers retain independent security baselines and emergency distrust powers (Mozilla, 2023c) [P4]; and that ETSI profiles aim to align with CA/B Forum practices so that qualified identity can coexist with WebPKI norms (ETSI, 2024a, 2025) [P2, P3]. Industry position papers further claim that implementing 2-QWAC avoids the need to add QTSP roots to TLS stores (Bitkom, 2020) [P3].

[CI-S1] QWACs integrate Website AuthN into EU Trust Schemes

Claim/Conclusion: QWACs standardise organisational-identity validation and support cross-border interoperability in the EU by anchoring issuance, profiles, and supervision in ETSI standards and by relying on the MS-TL/LoTL mutual-recognition regime. [P1, P2, P3, P4]

Data/Evidence: The amended Article 45 obliges user agents to recognise QWACs and display qualified identity. [P1] QTSPs operate under general policy requirements and supervisory controls defined by EN 319 401. (ETSI, 2024a) [P1, P3] Identity semantics for website certificates are harmonised in EN 319 412-4, including organisational names, jurisdictional data, and registration identifiers. (ETSI, 2023b) [P2] TS 119 411-5 v2.1.1 specifies browser-facing consumption outcomes for 1-QWAC and 2-QWAC, enabling consistent retrieval and verification of qualified identity by relying parties. (ETSI, 2025) [P2] The MS Trusted Lists and the EU List of the Trusted Lists provide the registry infrastructure that operationalises mutual recognition across MSs. [P4]

Warrant/Explanation: When identity fields, validation steps, and supervisory criteria are uniform and tied to a common legal registry, relying parties can process certificates predictably across jurisdictions, reducing cross-border friction relative to heterogeneous CA practices or vendor-specific policies. (ETSI, 2023b, 2024a) [P2, P3, P4] Browser-facing outcomes in TS 119 411-5 align certificate semantics with client behavior for both in-chain and side-binding models, supporting interoperable presentation and verification. (ETSI, 2025) [P2]

Backing/Authority: The standardisation-to-interoperability link is grounded in primary EU law that mandates recognition and sets the supervisory framework. [P1, P4] It is further supported by ETSI specifications that define provider policy requirements, identity profiles, and consumption outcomes. (ETSI, 2024a, 2023b, 2025) [P2, P3] Supporter analyses argue that an EU-wide trust-service framework yields more consistent cross-border handling and a clearer four-corner trust model. (Bitkom, 2020; Wazan et al., 2024) [P1, P2, P4]

Qualifier/Relativisation: Interoperability gains depend on close alignment with CA/B Forum and root-program practices to avoid divergence in acceptance and incident response. (Mozilla, 2023c) [P2] They also require unambiguous mapping from qualified identity data to browser UX and robust automation for discovery and status of qualified identity artefacts, especially for 2-QWAC. (ETSI, 2025) [P2] Consistent and effective supervisory practice across Member States remains essential. (ETSI, 2024a) [P3] The corpus provides limited independent, at-scale deployment evidence for seamless EU-wide operation under Article 45. (European Commission, 2021d) [P1]

Rebuttals/Counter: Critics contend that QWACs can fragment rather than unify practice by introducing optional out-of-chain identity bindings and governance anchored in MS-TL/LoTL, creating conflicts with browser root-store policies and incident-response baselines. (CCADB, 2020a; Mozilla, 2021, 2023c; Rescorla, 2022) [P2, P4] Program memoranda highlight additional discovery, verification, and revocation plumbing not covered by existing relying-party automation for the WebPKI. (CCADB, 2020a) [P2]

[CI-S2] **QWACs integrate with EU Cybersecurity Directives**

Claim/Conclusion: QWACs align with EU cybersecurity policy by providing a supervised, standardised organisational-identity control inside the eIDAS trust-services framework that user agents are required to recognise and surface, enabling cross-border, auditable attribution consistent with EU governance aims. [P1, P2]

Data/Evidence: Regulation (EU) 2024/1183 amends Article 45 to obligate user agents to recognise QWACs and to display the attested identity in

a user-friendly manner (ESD, 2023a). eIDAS supervision embeds organizational assurance through QTSP listing and periodic conformity assessment grounded in EN 319 401 and EN 319 411-1 (ETSI, 2024a, 2023b). EN 319 412-4 specifies organisational identity content, and TS 119 411-5 V2.1.1 defines validation and the two consumption patterns (1-QWAC and 2-QWAC) for relying parties (ETSI, 2023d, 2025). ENISA frames QWACs as instruments to promote trust, accountability, and transparency in the website authentication market (ENISA, 2016, 2017). Commission evaluations and policy materials present the rationale that harmonized, qualified identity assurance complements EU-level cybersecurity and market integrity objectives (European Commission, 2021d,b; Bitkom, 2020). [P1, P2]

Warrant/Explanation: If EU cybersecurity governance seeks verifiable attribution, supervised assurance, and interoperable enforcement across Member States, then a legally anchored identity artefact with harmonised semantics and supervisory audits can reduce ambiguity about controller identity for compliance and relying-party processes, while TS 119 411-5 ensures that transport-layer cryptography remains unaffected under the 2-QWAC model (European Commission, 2021b; ENISA, 2017; ETSI, 2025). [P1, P2, P3]

Backing/Authority: The legal duty and governance hooks are provided by Regulation (EU) 2024/1183 (Arts. 45 and 45a); operational baselines for QTSPs and identity semantics are specified in EN 319 401 and EN 319 411-1/412-4 (ETSI, 2024a, 2023b); consumption and interoperability guidance for QWACs are defined in TS 119 411-5 (ETSI, 2025); strategic positioning and intended policy fit are articulated in ENISA guidance and supportive industry and policy papers (ENISA, 2016, 2017; European Commission, 2021d; Bitkom, 2020). [P1, P2]

Qualifier/Relativisation: The contribution to cybersecurity integration is conditional: the Regulation is technology-neutral about encodings and does not mandate linkage to incident-response or CT ecosystems; benefits depend on uniform implementation and adoption, stable user-agent UI, and effective discovery and revocation plumbing under 2-QWAC (ETSI, 2025). The corpus provides policy rationales but little outcome-level evidence that QWACs measurably improve NIS-related compliance or incident rates relative to existing controls (Laurie et al., 2021;

Certificate Transparency, 2020). [P1, P2, P4]

Rebuttals/Counter: Browser and civil-society critiques argue that mandatory recognition and identity display may conflict with root-store governance, add integration and audit costs, and provide limited practical uplift given the historical weakness of identity-salience UI; they contend that modern DV+CT practices and agile browser policy already deliver accountability without statutory UI mandates (Mozilla, 2021; Mozilla et al., 2023; Internet Society et al., 2023; CCADB, 2020a; Felt et al., 2016; Laurie et al., 2021). [P3, P4]

[CI-S3] **QWACs strengthen Accountability and Transparency**

Claim/Conclusion: QWACs can advance GDPR-aligned transparency and accountability by making the legally responsible website operator unambiguously identifiable at the point of interaction and by obliging user agents to surface that identity in a user-friendly manner. [P1, P2]

Data/Evidence: Article 45 as amended requires user agents to recognise QWACs and display the attested identity [P1]. ETSI EN 319 412-4 harmonises identity attribute mapping to a legal entity (e.g., organisationName, registered jurisdiction, registration number) [P2] (ETSI, 2023d). EN 319 401 sets general policy requirements for trust service providers that underpin identity assurance [P2] (ETSI, 2024a). ETSI TS 119 411-5 v2.1.1 specifies relying-party consumption outcomes for 1-QWAC and 2-QWAC so that identity can be consistently retrieved and presented [P2] (ETSI, 2025). ENISA guidance and the EC CEF ntQWAC pilot describe QWACs as instruments to promote consumer trust and accountability by verifying controllers [P1, P2] (ENISA, 2016, 2017; European Commission, 2020a). Practitioner-advocacy materials explicitly connect qualified identity disclosure to GDPR transparency and accountability [P1] (Bundesdruckerei, 2022; Bailey, 2022).

Warrant/Explanation: When controllers are clearly and verifiably identified at the moment of collection or interaction, information duties and accountability chains can be attributed to a specific legal person, enabling cross-border enforcement and redress [P1] (ENISA, 2017). Harmonized identity semantics and profiles reduce ambiguity across Member States and support automated due diligence by enterprises and intermediaries [P2] (ETSI, 2023b, 2025).

Backing/Authority: The obligation to recognise and display qualified identity is in Regulation (EU) 2024/1183 [P1]. Identity content and assurance controls are standardized in EN 319 412-4 and EN 319 401, with deployment and consumption outcomes in TS 119 411-5 [P2] (ETSI, 2023b, 2024a, 2025). ENISA and EC materials frame QWACs as mechanisms for transparency about who operates a site [P1] (ENISA, 2016, 2017; European Commission, 2020a).

Qualifier/Relativisation: The GDPR link is indirect: qualified identity display alone does not ensure lawfulness, fairness, or purpose limitation, and its effect depends on correct binding, comprehensible UI, broad adoption, and reliable discovery and status handling, especially for 2-QWAC [P3, P4] (ETSI, 2025; Mozilla, 2023a). Empirical usable-security work shows limited behavioral impact from identity indicators, and we found no controlled studies demonstrating improved GDPR outcomes attributable to QWAC display at the population scale [P3] (Felt et al., 2016; Thompson et al., 2019; Feisty Duck, 2019).

Rebuttals/Counter: Critics argue that UI identity salience historically failed to change user behavior (EV de-emphasis) and that DV+CT+short lifetimes already provide effective, lower-complexity safeguards without mandating identity display [P3] (Felt et al., 2016; Thompson et al., 2019; Feisty Duck, 2019; Certificate Transparency, 2020; Mozilla, 2023c). They further note 2-QWAC introduces new verification plumbing and potential privacy or UX risks, and raises interoperability issues relative to existing relying-party automation [P4] (Mozilla, 2022, 2023a; CCADB, 2020a; Rescorla, 2022; Muffett, 2023a).

[CI-C1] **QWACs create Fragmented Trust Ecosystems**

Claim/Conclusion: QWAC adoption risks fragmenting the WebPKI by introducing parallel recognition and verification paths tied to MS-TL/LoTL governance and optional identity bindings outside the TLS certificate chain, leading to inconsistent validation, display, and incident-response behavior across user agents and jurisdictions. [P1, P2, P3, P4]

Data/Evidence: Regulation (EU) 2024/1183 obliges user agents to recognize QWACs and to surface attested identity in a user-friendly manner. [P1] ETSI TS 119 411-5 v2.1.1 defines two consumption models

(1-QWAC, 2-QWAC), with 2-QWAC leaving TLS transport semantics unchanged while adding implementation choices for discovery and binding (ETSI, 2025). [P2, P4] The eIDAS trust framework relies on Member State Trusted Lists aggregated via the List of Trusted Lists and supervision of QTSPs, a governance track distinct from browser root programs (European Commission, 2021a; ETSI, 2023e). [P3] The CCADB documents interoperability challenges when identity data are conveyed outside the TLS chain and must be discovered, verified, and status-checked through new client plumbing (CCADB, 2020b,a). [P2, P4] Browser program analyses and root store policy note misalignment risks between statutory recognition paths and root-store governance, which can yield divergent acceptance, UI, and incident-response outcomes (Mozilla, 2021, 2023c). [P3, P4] Technical analyses emphasize that QWACs do not change TLS protocol semantics; therefore, any benefit depends on correct and consistent client consumption logic that varies across implementations (Rescorla, 2022; Helme, 2023). [P2, P4] Early vendor materials exploring eIDAS 2 UI options illustrate multiple possible display patterns, increasing variability unless standardized (Entrust Corporation, 2021). [P4]

Warrant/Explanation: If identity assurance is verified and governed under statutory trust lists and may be conveyed via a qualified artifact outside the certificate chain relied on by root-store policy and CT-centric oversight, then relying parties must implement additional, non-standard discovery, verification, status, and UI code paths; under heterogeneous implementations and supervisory practices, this predictably produces inconsistent user experience, policy enforcement, and incident handling across the ecosystem (CCADB, 2020b; Mozilla, 2021). [P2, P3, P4]

Backing/Authority: CCADB interoperability notes and program discussions document technical and governance deltas for out-of-chain identity bindings (CCADB, 2020b,a); Mozilla policy and root-store materials analyze misalignment risks between statutory recognition and browser-governed trust stores (Mozilla, 2021, 2023c); research on non-governmental governance of the WebPKI explains how duplicative trust paths can fragment operational practice (Grindal et al., 2025). [P3, P4]

Qualifier/Relativisation: The fragmentation risk is mitigable: the legal text is technology-agnostic and does not prescribe a specific UI; ETSI 119

411-5 v2.1.1 aims for compatibility with existing practice and preserves TLS semantics via the 2-QWAC model, and vendors could standardize discovery, status, and display to reduce variance (ETSI, 2025). Our corpus contains no post-adoption outcome studies that empirically demonstrate pervasive fragmentation under the amended regulation. [P1, P2]

Rebuttals/Counter: Supporters argue that qualified profiles are designed for compatibility with the WebPKI, that 2-QWAC intentionally leaves TLS unchanged, and that harmonized identity semantics reduce UX variance; these claims draw mainly on standards-body and proponent materials (ENISA, 2017; Entschew et al., 2022; Bitkom, 2020; Bundesdruckerei, 2022). [P2, P3]

[CI-C2] **QWACs undermine Technological Neutrality Principles**

Claim/Conclusion: A mandatory QWAC pathway risks violating the EU principle of technological neutrality by prescribing recognition and identity display of a specific legal-entity artefact in user agents, which can crowd out vendor-led alternatives such as DV+CT with short lifetimes.

Data/Evidence: Regulation (EU) 2024/1183 amends Article 45 to oblige user agents to recognise QWACs and to ensure that the attested identity is shown in a user-friendly manner, and it sets supervisory procedures relevant to distrust decisions [P2]. ETSI TS 119 411-5 v2.1.1 specifies the implementation and consumption outcomes for qualified website certificates but does not dictate a particular user interface [P3] (ETSI, 2025). Industry and civil society statements specifically warn that Article 45 privileges QWAC-based legal-entity signalling over browser-security practices and narrows design discretion [P1, P2] (Mozilla et al., 2023; Internet Society et al., 2023; Mozilla, 2021). In contrast, the prevailing browser security baseline emphasizes DV plus Certificate Transparency auditing and increasingly short certificate lifetimes as evidence-driven mitigations [P1, P4] (Laurie et al., 2021; Certificate Transparency, 2020).

Warrant/Explanation: When law compels user agents to recognise and surface a particular identity artefact and gates distrust through supervisory processes, user-interface and enforcement choices become legally

coupled to that artefact, reducing the scope for evidence-driven alternatives such as DV+CT-only flows, tight lifetimes, or different identity presentations [P1, P2, P4] (Mozilla, 2021; Internet Society et al., 2023; Mozilla et al., 2023).

Backing/Authority: The binding legal obligation to recognise and display QWAC identities is codified at EU level [P2]. The ETSI documents delineate the qualified-certificate family and its implementation profile without prescribing UI specifics, reinforcing that the artifact is certificate-class specific rather than UI-agnostic policy [P3] (ETSI, 2025, 2024a, 2023e). Independent critiques by browser vendors and civil society groups articulate neutrality and discretion concerns in detail [P2] (Mozilla, 2021; Mozilla et al., 2023; Internet Society et al., 2023).

Qualifier/Relativisation: The Regulation is technology-agnostic in form and does not mandate a particular encoding or UI; ETSI TS 119 411-5 can be implemented with offline validation, and nothing in Article 45 forbids CT policies or short certificate lifetimes [P3, P4] (ETSI, 2025; Laurie et al., 2021).

Rebuttals/Counter: Supporters argue neutrality is preserved because the law recognizes a class of qualified certificates rather than a specific technology component, ETSI profiles aim for compatibility with existing WebPKI requirements, and UI choices remain with implementers [P3] (ETSI, 2025, 2023e, 2024a; Entschew et al., 2022; Martius et al., 2024). Proponents also contend that clear legal-entity signaling can enhance user trust when surfaced well, without defining how [P3] (Entschew et al., 2022; European Commission, 2021a).

[CI-C3] **QWACs increase Complexity/Costs for Website Operators**

Claim/Conclusion: QWAC qualification entails additional complexity and non-trivial costs, especially for SMEs, because issuers must satisfy ETSI qualification alongside existing WebPKI BR and root-program obligations, and relying parties may need new verification and display paths beyond BR-only WebPKI, notably for 2-QWAC [P1, P2, P3].

Data/Evidence: Only QTSPs may issue QWACs; QTSPs are supervised and listed on MS-TLs aggregated by the LoTL (European Commission,

2021a, 2023a) [P1]. QTSP qualification requires periodic conformity assessment against EN 319 401 and EN 319 411-1/-2. TS 119 411-5 v2.1.1 defines 1-QWAC and 2-QWAC consumption outcomes and profiles, with 2-QWAC using an identity binding outside the TLS chain that introduces issuer and client implementation work (Rescorla, 2022) [P2]. CCADB documents extra discovery, verification, and revocation plumbing for non-chain bindings beyond existing TLS chain automation (CCADB, 2020b). Mozilla program and policy materials describe mandatory recognition and identity-surfacing duties and the associated engineering and UI efforts (Mozilla, 2021, 2022, 2023b). Public WebPKI controls continue in parallel (CA/B Forum BR, root-store policy); therefore, ETSI qualification is additive rather than a replacement (CA/B Forum, 2025; Mozilla, 2023c) [P3].

Warrant/Explanation: If issuers must maintain partially overlapping assurance regimes (ETSI/CCAB) and clients must implement additional verification, status, and display paths for qualified identity (especially for 2-QWAC), then fixed and variable costs rise relative to BR-only issuance and consumption; those costs weigh more on smaller providers that cannot easily amortize new tooling (CCADB, 2020b; Rescorla, 2022; Mozilla, 2021) [P2, P3].

Backing/Authority: Binding legal texts encode obligations and supervision; ETSI EN 319 401/411 and TS 119 411-5 define qualification requirements, identity content, and consumption outcomes [P1, P2]; CA/B Forum BR and root policies define parallel WebPKI governance (CA/B Forum, 2025; Mozilla, 2023c); CCADB memoranda identify interoperability and automation gaps for non-chain bindings (CCADB, 2020b); browser-program materials describe engineering and UI obligations in practice (Mozilla, 2021, 2023b) [P3].

Qualifier/Relativisation: Burden is context-dependent and may be mitigated: TS 119 411-5 aligns qualified profiles with BR practice and preserves TLS semantics via 2-QWAC (ETSI, 2025); ETSI TR 119 411-5 provides coexistence guidance that can reduce integration friction (ETSI, 2023f); established QTSPs may realize economies of scale; automation for binding discovery, verification, and status could lower client costs; Article 45 obliges user-agent recognition and identity display but does not mandate site-side adoption of QWACs. The corpus

contains no quantitative, neutral total cost of ownership (TCO) or SME-specific impact studies on QWAC deployment. [P1, P2, P3, P4]

Rebuttals/Counter: Supporters argue that qualification largely reuses existing controls, that 2-QWAC avoids perturbing TLS, and that standards harmonisation plus supervision yield efficiency and accountability over time (Entschew and van Brouwershaven, 2024; Bitkom, 2020); these claims are mainly advanced in standards and pro-QWAC sources and are not accompanied by independent cost data. [P2, P4]

[CI-C4] **QWACs underperform Compared to Existing Measures**

Claim/Conclusion: Within the shared premises that QWACs do not change TLS cryptography and that the WebPKI already operates under mature, programmatic controls [P2, P3], the combination of DV under the BR, comprehensive CT logging and monitoring, short certificate lifetimes with ACME-based automation, and agile browser incident response provides security benefits comparable to those sought via QWACs at materially lower technical and governance complexity.

Data/Evidence: Empirical sources report that: (i) EV identity indicators were removed or demoted with no measurable improvement in user protection, weakening claims that additional positive identity salience reduces phishing (Thompson et al., 2019; Felt et al., 2016; Feisty Duck, 2019; Biddle et al., 2009) [P3]; (ii) CT is the de facto mechanism for transparent issuance oversight and rapid mis-issuance detection at scale (Certificate Transparency, 2020; Laurie et al., 2021) [P3]; (iii) short validity periods together with automated reissuance reduce exposure from key compromise and mitigate reliance on brittle user-facing revocation paths (Aas et al., 2019; Mozilla, 2023c) [P3]; (iv) root-store programs demonstrate fast, unilateral incident response that limits blast radius from CA failures (Wilson, 2015) [P3]; and (v) proposed out-of-chain bindings for 2-QWAC introduce extra discovery, verification, and status plumbing that existing RP automation does not cover, whereas DV+CT+short-lifetimes reuse deployed hooks (CCADB, 2020a; CA/B Forum, 2020; Mozilla, 2021; Helme, 2023) [P2, P3].

Warrant/Explanation: If dominant risks for website authentication are operational rather than cryptographic (mis-issuance, stale status, long life-

times, social engineering), then controls that maximize issuance transparency via CT, compress exposure windows (short lifetimes with automated renewal), and enable swift program enforcement (root-store actions) more efficiently reduce successful impersonation and limit CA failure impact than adding a separate, legally governed identity artifact that creates new verification paths and potential failure modes (Rescorla, 2022; Certificate Transparency, 2020; Mozilla, 2023c; Drury and Meyer, 2019; Amann et al., 2017) [P3].

Backing/Authority: Technical analysis holds that QWACs do not strengthen TLS semantics and add verification surfaces (Rescorla, 2022) [P2]; CT design and deployment materials establish transparency as the preferred issuance oversight control (Certificate Transparency, 2020; Laurie et al., 2021) [P3]; UI studies and field-scale evidence show limited or no protective effect from positive identity indicators (Thompson et al., 2019; Felt et al., 2016; Biddle et al., 2009) [P3]; incident-response practice in root programs demonstrates effective containment of authority failures (Wilson, 2015) [P3]; and CCADB and CA/B Forum memoranda document interoperability and automation gaps for non-chain QWAC bindings (CCADB, 2020a; CA/B Forum, 2020) [P2].

Qualifier/Relativisation: DV+CT+short-lifetimes do not provide public-law accountability or harmonized legal-entity semantics, which are central QWAC objectives (ENISA, 2016, 2017) [P1, P2]; their effectiveness depends on vigilant monitoring, CT log ecosystem health, and strict root-program enforcement (Certificate Transparency, 2020; Mozilla, 2023c) [P3]; and the corpus contains no controlled, post-amendment field studies that directly compare user-harm outcomes between DV+CT+short-lifetimes and QWAC-mediated identity presentation (Tehrani et al., 2024; Mozilla, 2021) [P4].

Rebuttals/Counter: Supporters argue that QWACs add supervised, enforceable organizational attribution and a legal duty to surface identity, with ETSI profiles specifying identity semantics and consumption models that preserve TLS while enabling consistent presentation in regulated or cross-border contexts (ETSI, 2023b, 2025; ENISA, 2017; Bitkom, 2020) [P1, P2]. They further contend that offline validation and careful engineering can bound privacy and latency, avoiding per-visit network leakage beyond modern DV practices (ETSI, 2025; ENISA, 2017) [P2].

Appendix C. Argument-linked Threat Model

Table C.1: QWAC threat model (WebPKI/TLS lens) under eIDAS 2.0 and ETSI TS 119 411-5 v2.1.1.

ID	Threat	Preconditions	Impact (Stakeholder)	Design levers
T1	Operator identity ambiguity enables fraud and weak attribution [ST-S1, ST-S3, CI-S3]	Look-alike domains and branding; weak domain-to-entity mapping; identity attributes non-unique, unstable, or not machine-readable; relying parties map free-text names into allowlists/policy	Phishing and fraud; slower incident triage and investigation; weaker accountability and recourse. (End Users; Website Operators; Browsers; Browser Vendors; QTSP; ETSI; NSB; MS; EC)	Minimal, stable, machine-readable identity set; include a unique registry identifier (e.g., BRN or LEI); integrate into enterprise anti-phishing and policy engines; avoid UI semantics implying “safe”
T2	Qualified UI becomes trust laundering and over-trust [ST-S2, ST-C1, ST-C2]	UI uses positive security cues; users infer “safe” from “identified”; attacker obtains QWAC for a misleading-but-legal entity; name/brand similarity exploited	Amplified social engineering; false confidence; reputational harm when “qualified” sites abuse trust. (End Users; Website Operators; Browsers; Browser Vendors; QTSP; ETSI; EC; MS)	Strict separation of identity vs connection security; identity shown in a neutral panel (not a lock/badge); mismatch- or policy-based warnings; usability testing before any prominent indicator; avoid “safe site” language
T3	Mis-issuance and identity-proofing failure, incl. shell entities [ST-S1, ST-C4, CI-S3]	Weak vetting and cross-checking; inconsistent registries across MS; fraudulent documents or shell entities; audit gaps; supervision and escalation gaps	Wrong legal identity displayed with qualified authority; fraud enablement; loss of trust in scheme and oversight. (End Users; Website Operators; Browsers; Browser Vendors; QTSP; CAB; Auditors; NAB; NSB; MS; EC)	Vetting controls and registry cross-checks; require stable unique identifiers; mandatory incident escalation and revocation SLAs; short(er) lifetimes; issuance transparency/monitoring (CT Logs or equivalent where applicable)
T4	Issuer key compromise or rogue/compelled issuance enables interception capability [GS-C1, ST-C4]	Issuer signing key compromise, coercion, or insider abuse; clients accept QTSP trust anchors for TLS (1-QWAC) and/or identity UI (2-QWAC); slow detection/distrust	1-QWAC enables TLS MITM; 2-QWAC enables false qualified identity; systemic trust failure and emergency distrust events. (End Users; Website Operators; Browsers; Browser Vendors; Root Store Operators; QTSP; NSB; MS-TL; LoTL; EC; MS)	HSMs and strict key ceremonies; separation of duties; continuous monitoring; short lifetimes; fast revocation and emergency distrust; prefer 2-QWAC for general web to keep TLS anchored in existing WebPKI controls

Continued on next page

ID	Threat	Preconditions	Impact (Stakeholder)	Design levers
T5	Trust-list poisoning or stale trust-list state [GS-C2, CI-C1]	MS-TL/LoTL signing compromise; malicious or erroneous QTSP entry; delayed removal propagation; client caching or rollback to stale lists	Malicious issuers remain trusted; delayed containment; inconsistent trust state causes outages and fragmentation. (Browsers; Browser Vendors; Root Store Operators; QTSP; MS-TL; LoTL; NSB; MS; EC)	Harden trust-list signing keys; transparent trust-list diffs and update feeds; independent monitoring; freshness and rollback protections; operational SLAs for urgent removals
T6	Governance mismatch and distrust latency (statutory process vs root-store incident response) [GS-S1-S3, GS-C2, GS-C3]	High-severity incident spanning statutory supervision and root-store policy; unclear authority boundaries; cross-border process friction; no single escalation channel	Slow containment and larger blast radius, or fragmentation and service disruption under unilateral action. (EC; MS; NSB; Root Store Operators; Browser Vendors; Browsers; QTSP; Website Operators; End Users; ETSI)	Single escalation channel (one PoC); codify an operational “security override” for active exploitation; align operational criteria with root-store practice where possible; keep baseline TLS policy unchanged before QWAC-specific logic
T7	Expanded TLS trust surface under 1-QWAC [CI-C1, GS-C3]	1-QWAC deployed as TLS server certificate; broad acceptance of additional QTSP trust anchors for TLS termination; heterogeneous clients and uneven baseline requirements	Larger blast radius from issuer failure; inconsistent reachability; increased systemic WebPKI risk. (End Users; Website Operators; Browsers; Browser Vendors; Root Store Operators; QTSP; EC; MS; CA/B; CT Logs)	Constrain 1-QWAC scope (if used); prefer 2-QWAC for general web; require baseline browser TLS requirements (e.g., CT Logs and revocation policy) before any added trust anchors; treat QWAC as additive identity check, not a TLS replacement
T8	2-QWAC TLS certificate binding suppression, downgrade, or staleness [ST-C2, CI-C1]	On-path attacker blocks binding retrieval; binding resource unavailable; caching/expiry mistakes; TLS certificate rotation; fail-open qualified UI	Qualified identity missing when intended or wrong identity shown; enterprise policy bypass. (End Users; Website Operators; Browsers; Browser Vendors; QTSP; ETSI)	Fail-closed for showing any “qualified” indicator; explicit “identity unavailable” state; strict binding validation (signature plus match to current TLS certificate); robust caching and freshness rules; conformance tests and vectors
T9	Revocation and status-check gaps across both planes [CI-C1, ST-C4]	OCSP/CRL/trust-list outages; soft-fail behavior; inconsistent semantics; 2-QWAC introduces additional artifacts to check and cache; delayed distrust propagation	Extended exposure after compromise or mis-issuance; inconsistent client behavior and UI state. (End Users; Website Operators; Browsers; Browser Vendors; Root Store Operators; QTSP; Revocation Registries (CRL, OCSP); MS-TL; LoTL; NSB)	Short-lived certs; stapled or offline status where feasible; consistent semantics for qualified-identity validity vs TLS validity; rapid distrust playbooks; monitoring for status failures

Continued on next page

ID	Threat	Preconditions	Impact (Stakeholder)	Design levers
T10	Privacy leakage from discovery, validation, and identity disclosure [ST-C3]	Per-visit online lookups (binding/status); third-party endpoints; mobile/WebView contexts; identity fields reveal more operator data than needed by default	Browsing metadata leakage and correlation; tracking; compliance and trust concerns. (End Users; Browsers; Browser Vendors; Website Operators; QTSP; Revocation Registries (CRL, OCSP); ETSI; EC)	Avoid per-visit “phone home”; caching and batching; stapled proofs where possible; minimize default identity disclosure; privacy review gates for any required network fetches
T11	New exploit surface from QWAC parsing and binding verification code [ST-C4]	New/immature validation code (ETSI profiles, trust-list parsing, binding verification); insufficient fuzzing and differential testing; ambiguous edge cases; divergent implementations	Validation bypass; spoofed identity UI; memory-safety vulnerabilities. (Browsers; Browser Vendors; QTSP; ETSI; Root Store Operators; MS-TL; LoTL)	Reuse hardened libraries; fuzzing and differential testing; minimize accepted profile surface; publish interoperability suites and test vectors for trust-list parsing, QWAC validation, and binding verification
T12	Operational complexity, cost, and misconfiguration risk [CI-C3, CI-C4]	Limited automation; partial client support; multiple artifacts/profiles; SME capacity constraints; monitoring and incident workflows not integrated	Outages and misconfigs; inconsistent identity display; low adoption; opportunity cost vs simpler TLS hygiene (DV, CT Logs, short lifetimes). (Website Operators; Browsers; Browser Vendors; QTSP; End Users; ETSI; EC; MS)	Automation and deployment playbooks; staged pilots in high-value contexts; success metrics (fraud outcomes, response latency, adoption); ensure identity failures degrade identity only and not TLS security